

**DISPLAY TECHNIQUES FOR PILOT INTERACTIONS
WITH INTELLIGENT AVIONICS:
A COGNITIVE APPROACH**

Marvin S. Cohen, Martin A. Tolcott, and James McIntyre

Prepared for:

**USAF Avionics Laboratory
Air Force Wright Aeronautics Laboratories
Attn: AFWAL/AAAT
Wright-Patterson Air Force Base, OH 45433-6543**

Contract No. F33615-86-C-1097

Submitted by:

**Decision Science Consortium, Inc.
7700 Leesburg Pike, Suite 421
Falls Church, Virginia 22043**

April 1987

TECHNICAL REPORT 87-6

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE APR 1987		2. REPORT TYPE		3. DATES COVERED 00-00-1987 to 00-00-1987	
4. TITLE AND SUBTITLE Display Techniques for Pilot Interactions with Intelligent Avionics: A Cognitive Approach				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Decision Science Consortium, Inc, 7700 Leesburg Pike, Suite 421, Falls Church, VA, 22043				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 163	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

ACKNOWLEDGEMENTS

This research was sponsored by the Avionics Laboratory, Air Force Wright Aeronautical Laboratories, Aeronautical Systems Division (AFSC), United States Air Force, Wright Patterson Air Force Base, Ohio, under Contract No. F33615-86-C-1097, as part of a six months Phase I effort in the Small Business Innovative Research (SBIR) program. We are grateful to Gurdial Saini, Lt. William Mallett, USAF, and Jerry Covert for their very constructive guidance during the project. We would also like to express appreciation to Capt. Steve Detro, Capt. Kevin Williams, and Maj. Joe Lutz, USAF, for their valuable contributions to our understanding of the pilot's point of view. Finally, appreciation is due to Theresa Mullin of DSC and Beth Adelson of Yale University for their stimulating technical contributions.

TABLE OF CONTENTS

<u>Section</u>	<u>Page</u>
1.0 INTRODUCTION	1
1.1 The Problem: Displays for Intelligent Systems	1
1.2 Objectives and Scope	3
1.3 Approach and Overview of the Report	5
2.0 COGNITIVE SCIENCE FOUNDATIONS FOR INTERFACE DESIGN METHODS	7
2.1 Levels of Cognitive Performance	7
2.2 Knowledge Representations	10
2.2.1 Schemas and scripts	10
2.2.2 Mental models	14
2.2.3 Analogical models and uncertainty	21
2.2.4 Hierarchical knowledge and the nature of expertise	28
2.2.5 Behavioral decision theory	32
2.3 Personalized and Prescriptive Decision Support: A Generalized Display Design Concept	38
3.0 PILOT KNOWLEDGE ELICITATION AND DESIGN OF DISPLAY CONCEPTS	49
3.1 Method	49
3.2 Uncertainty	51
3.3 Validity Checking of Data Sources	59
3.4 Hierarchical Knowledge Representation	65
4.0 CONCLUSIONS	70
4.1 Summary of Findings from Phase I	70
4.2 Future Directions	74
REFERENCES	77
APPENDIX	A-1

1.0 INTRODUCTION

1.1 The Problem: Displays for Intelligent Systems

Guidelines for the human factors engineering of the man-machine interface have traditionally focused on sensing and acting: i.e., display features and input devices that conform to human perceptual/motor capabilities and preferences. In recent years, however, artificial intelligence (AI) techniques have introduced a new class of systems with which humans are required to interact: systems which attempt to replicate, or improve on, human reasoning. As intelligent systems are proposed for an expanding sphere of operational roles, attention has begun to turn to machine-assisted thought, and to the manner in which computer-implemented storage and transformation of information can be optimally interfaced with human knowledge representations and processing strategies. Human-computer interface design has become *cognitive*. This report is intended as a contribution to the emerging application of cognitive science to human-computer interaction.

Nowhere is the challenge greater than in the design of pilot displays for intelligent avionics in high-performance combat aircraft. Near-future air warfare environments will be characterized by increasing aircraft velocities, by increasing sensor and weapon ranges, and by increasingly well-hidden threats on the ground and in the air. The result is both reduced response time for pilots and heightened uncertainty under which such responses must be made. Increasing automation of more routine system functions (such as aircraft control, target detection, tracking, and weapons control) has made cognitive activity, such as resolving uncertainty and balancing risks, a relatively more important and time-critical component of the pilot's task. The natural result has been increasing interest in the development of intelligent computerized support for high-level pilot decisions.

The interface problem for such systems is formidable. To work effectively, they must produce collaborative outputs that tap potential contributions of both human and computer within a period typically of a few seconds. In short, they must achieve a degree of cognitive integration of user and system that is virtually unheard of in other applications.

Traditional approaches to the human-computer interface (e.g., as summarized in Ramsey and Atwood, 1979; Engel and Granda, 1975) have not adequately addressed this problem. For example, principles for the design and formatting of displays are inadequate for the portrayal of abstract concepts, such as threat values and uncertainties regarding threat location and identity, on which tactical decisions (whether human or machine) must be based. Similarly, traditional guidelines for data entry are largely irrelevant for ensuring effective utilization of on-the-spot insights by users in a real-time process. Artificial intelligence contributions to the user-computer interface have focused for the most part on input-output tools (e.g., spatial data management, natural language understanding, voice I/O), rather than the effective use of those tools in collaborative human-computer problem-solving. Even work in the expert systems area (e.g., on explanation and mixed-initiative dialogues) has emphasized an essentially passive role for users, as initiators of queries, recipients of answers and explanations, and providers of raw, undigested data. One result has been the prevalent assumption that successful real-time tactical systems must entrust their duties almost wholly to the computer and leave little or no opportunity for human contributions. Collaborative aids that interweave human and computer reasoning and decision processes have evolved (if at all) by trial and error.

Efforts to develop a truly cognitive approach to interface design are, as yet, only incipient (cf., Norman and Draper, 1986). The enterprise is difficult for two reasons (at least). First, because cognitive science is itself not yet a mature discipline. A variety of models of human knowledge representation have been proposed, which differ in basic units (e.g., rules, objects, activities), in the processes that manipulate those units, and in the psychological functions they are thought to serve. Second, because the main focus is toward theory rather than application, the implications (if any) of a particular cognitive theory for the design of an interactive interface are often far from obvious. Research on the "application" of cognitive theories, therefore, is not a simple matter of converting first principles into engineering diagrams. It must itself proceed in a tentative, hypothesis-testing mode. First, concepts from basic research must be selected based on their apparent relevance to the problem domain and empirical plausibility; then, the

implications of these concepts for display design must be made explicit; finally, the displays based on these concepts must be carefully evaluated. The results, in turn, might provide valuable feedback and sharpened focus for basic research.

Our hypothesis is that recent work in cognitive science can provide the underpinnings for a new methodology of interface design for real-time interactive aids. Specifically, that methodology is based on insights from (a) work on knowledge representation and (b) research on psychological decision theory. These sources are complementary. Displays which represent information in accordance with users' own internal representations should be more readily utilized, should be understood more quickly and accurately, and should provide a more effective context for eliciting on-the-spot user knowledge. On the other hand, human knowledge representations and information processing strategies are imperfect; the literature on psychological decision behavior reveals a number of ways in which preferred methods for reasoning may lead to biases or fallacies. Our aim, then, is to articulate a design methodology which emphasizes both compatibility with user-preferred methods for representing and using knowledge, and techniques for avoiding the biases to which those methods ordinarily lead.

1.2 Objectives and Scope

The research reported here is the product of a 6-months Phase I effort in the Small Business Innovative Research (SBIR) program. The objectives were to:

- a) examine relevant theories and concepts from research on knowledge representation, behavioral decision making, and decision aiding,
- b) develop a methodology for generating display design concepts for pilot interaction with intelligent systems, based upon those theories,
- c) use the methodology to develop experimental display design concepts, and

- d) conduct preliminary feasibility tests of those concepts.

The initial application context involved an air-to-ground strike mission. In order to reach a target deep within enemy territory, an aircraft must avoid or defeat a variety of surface threats whose identity, location, and/or capabilities may be wholly or partly unknown. Information may be obtained during the flight itself from on-board sensors or radio messages from air or ground stations which in some cases can help identify new threats, resolve the uncertainties in prior intelligence, and help pilots select an adaptive response (e.g., a revised route). Several overlapping and interrelated topics were of specific initial concern to us within this context:

- o *dynamic displays*, that is, displays that change as the mission progresses, as new threat information is received, or as computations modify conclusions about threat assessments or preferred routes and tactics;
- o *uncertainty*, how pilots think about it, how it affects their decisions, and how displays should be designed to represent it;
- o *hierarchically organized information*, i.e., how information should be aggregated so that displays are uncluttered and the pilot's attention is focused on the appropriate level of detail;
- o *explanations of system reasoning*, i.e., how to display in a clearly intelligible way the basis of inferences from incomplete and unreliable data and the reasons for recommended courses of action within the limited available response time.

Phase I specifically excluded consideration of air threats and air-to-air missions. Further, the principle focus was on in-flight pilot aids as opposed to prestrike ground planning. Nevertheless, certain aspects of the planning process had to be considered, specifically the role of intelligence information and uncertainty about threat location and type, in order to understand the impact of new information received during the mission.

Finally, the emphasis in this phase of the study was on the development of a methodology based upon the underlying relevant theory, and on display concepts that illustrated the application of this theory, rather than on the development of detailed prototype displays. Thus, in balancing the amount of effort to be devoted during Phase I on theory and method versus software development, the emphasis was on theory and method.

1.3 Approach and Overview of the Report

The approach consisted of several steps:

- a) A critical review was conducted of the research literature dealing with knowledge representation (especially mental models) and behavioral decision theory (especially the work on cognitive biases leading to errors in judgment) in order to identify relevant theoretical formulations for an in-flight pilot display design methodology.
- b) Structured interviews were held with three experienced Air Force pilots of tactical strike aircraft, in which they were led through a typical mission, new threat information was presented periodically, and they were asked how they thought about the situation as it developed, the uncertainties inherent in the situation assessment, and the choice of responses. Questions were designed to probe their ways of mentally organizing and representing information, potential biases in making decisions, and the type of displayed information and method of display that would most help them in handling uncertainty and reaching a timely decision.
- c) The design methodology was applied to data elicited from the pilots, and a series of preliminary pilot displays was developed. The preliminary displays were programmed on an IBM-PC/AT in a sequence keyed to a mission scenario. These displays conformed to the constraints imposed by mental model theory, while providing prescriptive guidance based on behavioral decision theory.

- d) The demonstration system was reviewed individually by the three pilots who had been interviewed initially. Ratings were solicited from the pilots regarding specific features on each display. The ratings were based on a 7-point scale from 1 (very good) through 4 (neutral) to 7 (very bad), and comments were solicited to explain the reasons underlying the ratings and to suggest improvements or alternative designs. During this review (which was tape-recorded), the research team attempted to further clarify the mental models and decision strategies underlying the pilots' responses.
- e) Finally a demonstration version of the final display concepts was developed and demonstrated at Wright-Patterson Air Force Base.

In Section 2.0 below we examine the relevant cognitive science literature and describe a methodology for the design of displays for intelligent systems. Section 3.0 then presents the results of applying that methodology to the preliminary design and evaluation of in-flight pilot displays. Finally, Section 4.0 summarizes the conclusions from Phase I and points toward future research.

2.0 COGNITIVE SCIENCE FOUNDATIONS FOR INTERFACE DESIGN METHODS

In this section we propose a theoretical basis for a methodology of cognitive interface design. As noted in Section 1.1, that methodology has a dual basis: (1) displaying information in a way that is compatible with a decision maker's preferred method of representing knowledge and solving problems; while (2) providing protective devices to guard against associated biases. Thus, the two major areas of cognitive science research of concern to us are models of human knowledge representation and reasoning, and research on errors in judgment and decision making. This by no means, therefore, purports to be a complete review: there is considerable additional cognitive research literature with an important bearing on human-computer interaction. Rather, we focus here on work which, on the one hand, has been relatively neglected in the context of system design, and which, on the other hand, has been the major source of insights for the design methodology which we propose.

We argue that these two research traditions are complementary and can shed light on one another both at a theoretical level and in their application to design. Sections 2.1 through 2.3 examine this literature, while Section 2.4 extracts their implications for display design.

2.1 Levels of Cognitive Performance

Rasmussen (1983; 1986) has introduced a classification of levels of human performance which will serve as a useful starting point for the knowledge representation concepts to be developed in the next section. As shown in Figure 2-1, Rasmussen distinguishes performance which is skill-based, rule-based, and knowledge-based.

Skill-based behavior involves smooth, automated, highly integrated patterns of behavior in which the body typically acts as a "multivariable continuous control system synchronizing movements with the behavior of the environment." Sensory inputs serve two functions at this level: as "signs" which trigger appropriate behavioral patterns (e.g., an incoming missile elicits the response pattern of taking evasive action); and as "signals" which modulate and control an already activated pattern (e.g., observation of the distance

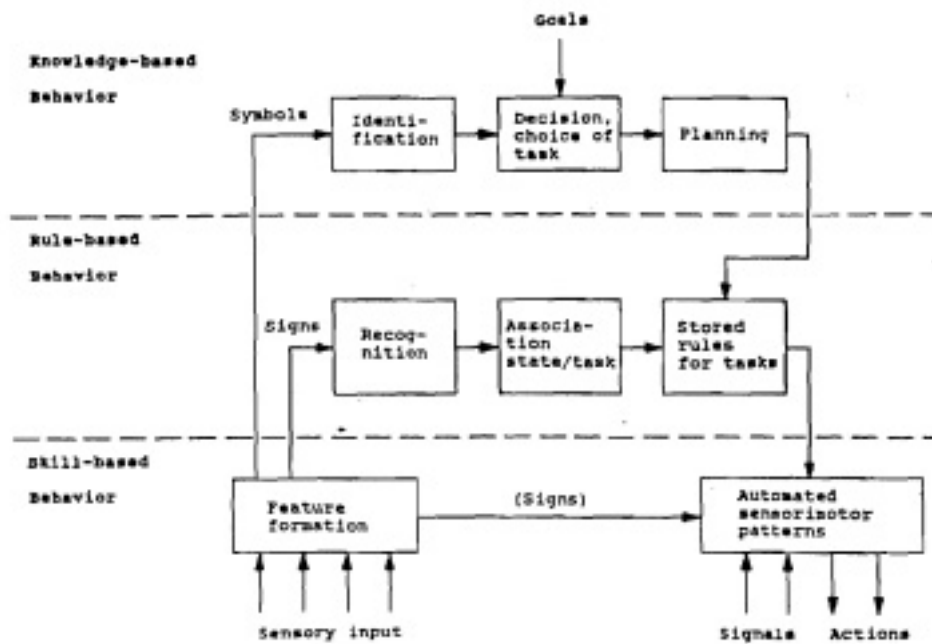


Figure 2-1: Rasmussen's Framework for Cognitive Performance

and angle of approach of the missile). Skill-based behavior is not typically a matter of simple feedback control. Rather, it depends on a flexible and dynamic internal model of the environment, which is continually updated by signals from the environment, which permits the individual to anticipate likely environmental perturbations, and which integrates activities into a single, smooth sequence. Pilots engaged in an evasive maneuver, for example, may have an instant three-dimensional mental "picture" of the entire pattern to be executed by the aircraft and an automated unconscious set of behavioral routines for carrying it out.

At the next higher level of performance, rule-based behavior is consciously controlled by a stored rule or procedure. Such a rule may have been acquired by direct experience or it may have been learned from other people by instruction. For example, a pilot may discover the appropriate distance and altitude for avoiding detection by a particular enemy missile site through his own experience (e.g., of being illuminated by its tracking radar at certain locations and not at others); or he may have been briefed on what to do in the vicinity of such a threat during mission planning. Rule-based behavior is goal-oriented only in a limited sense: behavior is governed by rules that were successful in previous performance (one's own or others'). But the goal remains implicit in the use of the rule; there is no explicit reasoning or problem-solving to discover the best way to achieve the goal. Individuals may acquire a large store of rules at this level which enable them to respond adaptively to relatively familiar or expected situations.

The next level of performance, knowledge-based behavior, is relevant when the situation is unfamiliar. If no rules are available for achieving the goal, the individual must draw upon a deeper understanding of the causal relationships in the environment which determine the conditions under which his goal can and cannot be achieved. He must construct a "mental model" of the situation in which alternative courses of action and alternative outcomes can be simulated. For example, a pilot confronted by conflicting information about the classification or location of a threat may utilize his understanding of the strengths and weaknesses of the different sources of information under various conditions to resolve the conflict. Confronted by an unusual configuration of unexpected threats on his flight path, he may utilize his

knowledge of threat capabilities and tactics to mentally "simulate" alternative routes. At this level, environmental inputs no longer function as "signs" which are associated with prelearned procedures, but as "symbols" which provide evidence for functional properties and causal relationships.

The development of cognitive capabilities often involves transfer of control from a higher to a lower level. Thus, an initial stage of rule-following (e.g., relying on instructor and textbook in the operation of a flight simulator) will be replaced after a period of direct practice by a more automated and intuitive mode of operation. Similarly, basic knowledge of the causal and functional properties of a domain (e.g., characteristics of weapons and threats) will be replaced after experience with more stereotyped rule-based reactions. Nevertheless, higher-level control may occasionally intervene in the execution of a well-practiced capability. Rules may control the sequencing of skilled routines or impose constraints on how the skill is executed (e.g., "the incoming missile is a very fast one, so execute the evasive pattern quickly.") Similarly, when rules prove inadequate for performance in novel situations (e.g., conflicting evidence or unexpected threats), it is necessary to ascend to higher level knowledge-based reasoning in order to determine what to do and, perhaps, to generate appropriate rules.

2.2 Knowledge Representations

Improved understanding of pilot cognitive processes can be obtained by going beyond Rasmussen's scheme to a more detailed consideration of the knowledge representations required to implement it: first, by introducing a more active and hierarchical representation of stereotypical information at the "rule-based" level; and second, by examining constraints on performance derived from the nature and function of mental models.

2.2.1 Schemas and scripts. In his discussion of performance in familiar situations, Rasmussen implies that knowledge of this type is composed of small unrelated units (rules) which are activated in a stimulus-driven, or "bottom up," fashion. An example of this type of control is a standard production system which contains a large number of rules of the form "If <situation> then <action>." When the conditions specified in the antecedent of the rule are

satisfied, the action described in the consequent is performed. That action may create conditions which cause other rules to fire, and so on.

A large body of cognitive science research suggests, however, that human performance even in familiar, stereotypical situations involves more highly structured types of knowledge and more active, "top-down" processing than are found in the standard production system. The notion of a schema (or frame) provides a convenient means of representing knowledge of this type. Schemas are data structures corresponding to familiar types of objects, situations, events, sequences of events, actions, or sequences of actions (Rumelhart and Norman, 1985).

Three features of schema-based representations are central:

- (1) Schemas have slots or variables which specify which types of information it is appropriate to seek about a particular type of thing. For example, a pilot's schema for a surface-to-air missile site might include slots for radar range, radar altitude, missile range, missile effectiveness, local terrain features, etc. In some cases slots may have default values, i.e., values which are expected or assumed to be correct until evidence to the contrary is obtained. For example, pilots may cautiously assume that a missile of unknown type has maximum capability until they learn otherwise.
- (2) Schemas represent knowledge at multiple levels, and these levels are hierarchically organized both in terms of "is-a-part-of" relationships and in terms of "is-a-kind-of" relationships. Schemas typically include other schemas as parts: e.g., the SAM site schema may include sub-schemas for each major component of the site (radar, missile, terrain). In addition, schemas may exist for types of objects, events, etc. at varying levels of generality; for example, there may be a general schema for weapons, a more specific schema for anti-air weapons, a still more specific schema for surface-to-air missiles, and a schema for a specific type of surface-to-air missile (e.g., SA-2). Each schema inherits slots and default values from schemas above it in the generalization hierarchy. Both types

of hierarchical organization provide powerful tools for generating expectations and guiding the collection of new information.

- (3) Finally, schemas are active processors of information, rather than static repositories of facts. The ensemble of schemas embodying a person's knowledge works together to make sense of incoming data and to guide action. Each schema continually assesses its own applicability to the current situation, determines what further information should be sought or expected, and forwards relevant findings to other schemas (cf., Minsky, 1979). Schemas thus bridge the traditional gap between static "declarative" representations and active "procedural" representations (Winograd, 1972).

The pattern of communication among schemas may be determined jointly by the task and by their hierarchical organization. In inference tasks like diagnosis or classification, a relatively generic schema (e.g., an antiair missile site) may respond to available evidence by deciding that it applies to a situation and then activate schemas for specific subtypes (e.g., SA-2, SA-8, SA-9 ...); the subtypes then compete to determine which of them applies; and so on down a hierarchical tree (cf., Chandrasekaran, 1983). The task of planning may also involve increasingly detailed specification of schemas. But planning may also involve the schema for part of an activity activating the schema for the whole, which in turn activates schemas for other (subsequent) parts of the activity. Schemas for activities thus not only support the process of recognizing or interpreting what is going on, but once the context has been recognized, determine what actions an agent should take within it (Galambos, Abelson and Black, 1986).

A specific type of schema, suited for representing knowledge about familiar activities, is the script (Schank and Abelson, 1977). A script contains slots whose values specify the objects ("props") and persons ("roles") which participate in the activity, its entry conditions, and results, as well as the sequence of scenes which constitutes the activity. For example, Figure 2-2 is a hypothetical pilot's script for an offensive counter-air (OCA) mission, and contains slots for the aircraft, target, IP, way-points, and mission components (or scenes). Such a script encapsulates the pilot's prestored

<u>SCRIPT</u>	OCA MISSION	<u>RESULTS</u>	DAMAGE TARGET RETURN
<u>ROLES</u>	PILOT ...	<u>ENTRY CONDITIONS</u>	ATO DIRECTIVE ...
<u>PROPS</u>	AIRCRAFT, ORDNANCE, JAMMING GEAR, TARGET, IP, WAYPOINTS		
<u>PREPARATORY EVENTS</u>	<u>ENABLEMENT EVENTS</u>	<u>PRECONDITION EVENTS</u>	<u>SIDE-CONDITION EVENTS</u>
PLANNING, EQUIPPING A/C ...	CROSS FEBA	GET TO TARGET BY TOT - ...	AVOIDING THREATS ON INGRESS IF AVOIDANCE IS POSSIBLE, AVOID; ELSE, USE CHAFF OR TRY TO KILL. IF SA-2 ... IF SA-6 ...
<u>ACTION EVENTS</u>	<u>POST-ACTION EVENTS</u>	<u>DISENABLEMENT EVENTS</u>	<u>TRANSITION EVENTS</u>
PROP ORDNANCE ON TARGET	AVOID THREATS ON EGRESS	RETURN TO BASE	BRIEF RESULTS ...

Figure 2-2: An Illustrative Script Representation for an OCA Mission

knowledge about OCA missions and guides his expectations and actions as he proceeds. (In a real script, scenes would be specified in much greater detail.)

Both types of hierarchical organization (is-a-part-of and is-a-kind-of) are relevant in script-based representations. First, scripts may be hierarchically composed of scenes (e.g., ingress, attack, egress) which are themselves scripts or which contain scripts as parts. Second, in recent work Schank (1982) has described more general schemas (called MOPS) of which scripts are instances. Planning may involve activation of relevant generic schemas followed by a process of filling in details until a specific script is constructed (cf., Stefik, 1981). For example, pilots may have generic schemas for strike missions which determine some of the features of OCA-mission scripts. On an even more general level, a schema called UM-Performance provides an abstract characterization which applies to OCA-missions as well as to any other "performance." According to Schank, this schema contains eight universal scenes: Preparatory (things done prior to entering a context), Enablement (entry into a context), Pre-Condition (things done prior to the main activity), Side-condition (tangential actions), Action (the main activity), Post-condition (tying up loose ends), Disablement (leave), and Transition (move on to new things). Figure 2-2 illustrates how scenes in the OCA mission script might be organized under these categories.

The concept of a script provides a considerably richer and more adequate representation of stereotypical performance than the notion of a rule. Scripts in fact provide a unifying context for other types of knowledge--both in the form of rules (at the most specific level within a scene) and in the form of other types of schemas (e.g., about threats, the aircraft, and terrain features), by showing where and how they become relevant in the course of a familiar activity (Leddo, Mullin, and Cohen, 1987).

2.2.2 Mental models. Scripts, by definition, provide no capability for dealing with novel or unexpected situations. When deviations from a script occur, or no familiar script can be found which adequately matches the given circumstances, knowledge-based reasoning must be invoked. Mental models may be employed to explain apparently anomalous events or to generate options that

overcome unanticipated obstacles. Given the severe time pressures constraining pilot performance, it is unlikely that they are able to make frequent effective use of mental models in this way. Nevertheless, the potential contribution of pilot knowledge at this level is great, as is the need for human-computer systems that can cooperatively adjust to unexpected circumstances. A major function of an intelligent avionics system may be to automate relatively routine or stereotypical tasks, and to alert pilots when high-level "managerial" or "troubleshooting" skills are required (cf., Moss, Reising, and Hudson, 1984). Thus, it will be worthwhile to explore theories about the way people naturally solve problems at this level.

What is a mental model? A variety of reviews and taxonomies of this concept now exist. Rouse and Morris (1986), building on Rasmussen (1979), provide a functional definition of mental models as "the mechanisms whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning and observed system states, and predictions of future system states." They then discriminate among mental models in different domains (e.g., problem solving in physics versus manual control) based on whether or not a person is aware of his or her manipulation of a mental model and the extent to which use of the model is a matter of choice as opposed to being dictated by the task. Young (1983) enumerates a variety of mental model mechanisms that have been proposed in the literature (e.g., analogy, device surrogate, mapping, problem space, grammar).

These discussions fail to provide a basis for the mental model concept which links the proposed properties of mental models (e.g., as reviewed by Young) to the functions they are meant to perform (e.g., as described by Rasmussen and Rouse and Morris). An understanding of that linkage is required in order to sort out and evaluate diverse definitions and theories. Moreover, Rouse and Morris do not distinguish between knowledge that is simply retrieved by means of a mental model and knowledge which is generated for the first time by such models. We would argue that a theory of mental models must in fact make this discrimination; that it should begin with the function of generating new knowledge and the constraints that function imposes on representational properties; and that the use of mental models to support stereotypical skill in familiar situations is derived from this more basic function.

In his classic book, *The Sciences of the Artificial*, Herbert Simon (1969) argues that "human problem solving, from the most blundering to the most insightful, involves nothing more than varying mixtures of trial and error and selectivity. The selectivity derives from various rules of thumb, or heuristics, that suggest which paths should be tried first and which leads are promising" (p. 97). Within our framework, Simon's "selectivity" corresponds to stereotypical, pre-existing knowledge, and "trial and error" corresponds to knowledge-based reasoning. As Simon notes, "the more difficult and novel the problem, the greater is likely to be the amount of trial and error required to find a solution" (p.95).

In a recent discussion, D.C. Dennett (1978) has argued that the prominence of generate-and-test (or "trial and error") mechanisms in AI programs is no accident, that any process in which genuinely new knowledge is created within a system must involve some version of variation (trial) and selection (error). First, since the knowledge is new, it must be underdetermined by the pre-existing design (i.e., knowledge) of the system; in other words, there must be a process of variation or option generation that is to some degree random or fortuitous. Variation by itself, however, can provide no more than a chance probability of improving on the old design. If the products of random variation are to produce new knowledge, there must be some process of selection which can reject variations on the basis of what is previously known.

Generate-and-test mechanisms operate at a variety of levels, which vary in the degree to which the processes of selection take place within the organism. In natural selection, variation in the genetic code may produce novel behavioral dispositions, perceptual/motor skills, etc. The selection process is not inside the organism at all; rather, it works through the differential survival of organisms in which such variations turn out to be environmentally adaptive. Generate-and-test mechanisms are also essential to learning at the individual level. Variations in an individual's behavior will be retained and reoccur when they produce environmental consequences which are perceived as rewarding (or prevent environmental events perceived as aversive). The selective events (the perception of reward or pain) are now inside the organism, but the environment still controls when they occur, by causally linking them to

behavioral variations. Knowledge-based reasoning carries the generate-and-test concept one step further, entirely internalizing the selective process: hypotheses are varied in an internal model of the environment, with selective retention of alternatives that prove successful *inside* that internal model.

The effectiveness of learning in creating genuinely new knowledge depends on two things: (a) the independence (i.e., "randomness") of the generation function (which is within the organism) with respect to the selection function (which is controlled by the environment); and (b) the selection process replicating to a reasonable approximation the effects of evolutionary selection; i.e., the pleasures and pains that shape an individual's behavior should be correlated to a degree with ultimate reproductive success and failure. That the latter is the case is largely ensured by the fact that the *capacity* for learning, and the particular events that serve as positive and negative consequences, have themselves evolved through natural selection. Similar conditions must apply for knowledge-based reasoning to be an effective method of creating new knowledge. The variation function within the organism must be independent (i.e., "random") with respect to the selection function since otherwise we have pre-existing (stereotypical) knowledge; and the selective function, also within the organism, must replicate (reasonably well) the selective action of the environment in learning--i.e., it must produce internal selective events in the same causally appropriate way as the environment does. At first glance, it is a mystery how this could be: either the knowledge that a particular variation produces a particular selective effect is already present in the organism (and so genuinely new knowledge is not produced) or the knowledge is not present (and intelligent selection is not possible).

At the knowledge-based level, then, variation and selection must at the same time be uncoupled within the organism (to achieve randomness of variations) and in another sense coupled (so that the selective function can "know" which variations are likely to be adaptive). Several important representational properties of mental models are suggested by the requirement that both of these conditions be simultaneously satisfied:

- o presence in the model of component(s) in which variations are represented (e.g., actions which are hypothesized to achieve an objective or states of affairs which are hypothesized to account for unexpected or anomalous events);
- o presence in the model of component(s) in which success or failure in achieving some selective criterion is represented (e.g., achievement of action goals; explanation of anomalous or unexpected events);
- o representation of relationships between variation component(s) and selection component(s) in such a way that when changes are made in the variation component(s), corresponding causally or logically appropriate changes occur in the selection component(s); and
- o absence of a pre-existing direct representation of these causal or logical relations (at any level of generality); for example, no explicit rules of the form "If <variation x> then <value y on selection criterion>" or "If <variation of type X> then <value of type Y on selection criterion>."

Detailed implementation of these properties might be accomplished in more than one way. However, we would argue that any successful implementation must involve certain common features: the notion that a mental model consists of multiple components, that pre-existing causal or logical knowledge about its own behavior is associated with each component, and that novel information about the adaptive adequacy of random variations is derived by "gluing" the components together and observing their interaction.

This is what a pilot does, for example, when in the face of unexpected threats he imagines an alternative route or an alternative set of tactics (e.g., ECM, chaff) and "plays out" the consequences of the option in his mind: what will each enemy unit think and do? what will he do in turn? etc. The components of the model are familiar (his own aircraft and its capabilities, the threats and their capabilities); but the configuration is novel. In order to evaluate an action, therefore, he must put the components together and internally "observe" their interaction.

A theory of mental models with these features has been developed by deKleer and Brown (1981). A mental model, on their view, consists, first, of a "device topology," i.e., a set of well-understood components, a set of well-understood "conduits" (connections by means of which components may causally affect one another), and a specification of which components are connected with which by conduits. Thus each component has a set of states it can be in (e.g., detected or undetected as states of own aircraft; detecting or not detecting as states of an enemy radar installation), and a set of rules determining how its state will change as a function of changes in the values of conduit attributes (e.g., ECM or chaff).

In terms of the previous section, we can understand deKleer and Brown's notion of a "device topology" as a system of schemas which represent knowledge about the properties and behavior of objects, and which send "messages" to one another representing cause-effect relationships and triggering state changes in the recipients. A key feature of this type of model is the *locality* of these cause-effect relationships; that is, rules for the behavior of any given component can only reference its own state and the attributes of the conduits connected to it, and can in no way refer to how the overall system is known or intended to function. For example, if the pilot already knows that if he adopts a certain tactic, he will not be detected, there is no point in utilizing a mental model. deKleer and Brown call this the "no-function-in-structure principle," and it represents the "uncoupling" which is essential for the model's ability to generate new knowledge from old knowledge. If mental models are to serve their purpose of generating predictions in novel circumstances, e.g., about the outcome of an option or the impact of a causal variable, they cannot rely on prior (stereotypical) knowledge of what is to be predicted.

deKleer and Brown's theory needs to be supplemented, however, by recognition that prior knowledge of a "non-local" sort (i.e., knowledge which goes beyond the information encapsulated in the separate object schemas) does play a crucial role in mental models, in at least two ways. First, as Simon (1969) noted, there is typically some selectivity in the generation of options for testing. We would argue more strongly that some selectivity must always be

present; otherwise, since there is an infinite space of potential solutions to be searched, adaptive possibilities would hardly ever be found. Prior (stereotypical) knowledge supports such selectivity: by narrowing the field within which options are generated for testing (e.g., there are some things the pilot already knows will not work), by providing components or building blocks for options which can be recombined in novel ways, or by bringing promising possibilities to mind (e.g., by analogy with some other situation the pilot has experienced or heard about. In the latter case, note that an analogy does not function as a general rule--e.g., "In all situations like x and y, do z"--but is more like a *hypothesis* that x and y are in fact similar: "In situation y, I did z and it worked. If situation x is like situation y, z may work here as well."). Only when prior knowledge fully determines the choice is the mental model not required.

Secondly, prior knowledge of a quite sophisticated sort is utilized in building the device topology. The pilot may never have encountered this specific configuration of threats, but he may be quite practiced at solving problems of this kind. He thus knows what components need to be included in the model (e.g., own aircraft, surface-to-air threats, terrain) and what parameters of each will be relevant. He "learns how to think" about such problems by developing abstract schemas or scripts for building appropriate mental models. Such abstract schemas and scripts may themselves be shaped by successes or failures in the real environment. Another possibility, which occurs both in science (Gentner and Gentner, 1983) and in ordinary reasoning (Lakoff and Johnson, 1980), is to construct mental models by metaphorically mapping objects and behaviors in one domain onto phenomena in another (e.g., conceiving electricity as a fluid, or an argument as a "war" between competing positions).

A device topology by itself is a static structure; it must be actively used if the required predictions are to be generated. The second major concept in deKleer and Brown's theory involves a process called "envisioning." Envisioning derives function from structure by a process of propagation whereby one starts with a single input state (e.g., an action option or candidate explanation), then examines the nearby components to observe its effects, examines the nearby components of those components, and so on. Envisioning

results in a dependency graph of causes and effects; e.g., if I do x, y happens; as a result, I do z, and w happens, and so on. In other words, envisioning converts a representation in terms of interacting objects (the device topology) into a representation of a set of temporally and causally related event schemas, which deKleer and Brown call the "causal model."

Although the basic idea of envisioning is quite simple, its application may in fact involve a quite difficult process of problem solving. The difficulty arises because the initial knowledge of device topology may be insufficient to determine the behavior of the system (e.g., if I do x, y might happen but z might also happen). When this is the case, deKleer and Brown propose that envisioning eliminates the ambiguity by making assumptions. Such assumptions may concern the existence of causally relevant but unobserved attributes, the temporal order of events, the satisfaction of rule conditions, or precise attribute values. Assumptions may have to be revised subsequently if actually observed events conflict with the events predicted by the model.

Once envisioning has produced a causal model (i.e., a predicted sequence of events), the model can be "run" to predict a specific event or outcome. Running is a relatively simple matter of activating a pre-existing schema. The main work of problem solving at the knowledge-based level has been accomplished by the processes in which the schema was created: i.e., constructing the device topology, generating an option, and envisioning its consequences.

2.2.3 Analogical models and uncertainty. We have argued that in order to support the function of generating new knowledge, mental models must involve some internal version of a generate-and-test process; and in order to implement the latter, mental models must be composed of well-understood components which are "glued" together in order to observe their interaction. We now consider an additional corollary of this argument, which has implications both for display design and for likely weaknesses in natural human methods of reasoning. Mental models which satisfy the above requirements belong to a class of models which may, somewhat loosely, be characterized as "analogical."

There is considerable discussion and debate in the research community regarding the nature of (and the need for) a distinction between "analog" and "propositional" representations (e.g., Pylyshin, 1979; Kosslyn, 1980; Rumelhart and Norman, 1985). Nevertheless, we would argue that a plausible and important distinction can be made, based on the requirement that mental models have the capability of generating new knowledge.

Shepard (1975) and Metzler and Shepard (1974) have summarized empirical evidence concerning the properties of mental images which appear to distinguish them from other internal representations. In particular, they mention:

- o a one-to-one correspondence between components of the representation and components of the situation which it represents (e.g., the image of a chair appears to have legs, a seat, a back);
- o a one-to-one correspondence in time of the states which the representation passes through and the states of the represented situation (e.g., in imagining the rotation of a three-dimensional object).

We would argue, based on our discussion in the previous sections, that both of these properties must characterize mental models. If such models are to adequately support an internal generate and test process, they must have components which correspond to components in the represented situation, and changes in the model components must causally mirror changes in the environment. In both respects, such models appear to differ from "propositional" representations, in which (a) there are syntactic elements (like "the" and "all") with no direct representational function, and (b) changes in state are more typically represented by large, abrupt shifts in the representation rather than a gradual transition through intermediate states.

Propositional representations can be developed which mimic the behavior of analogical models, i.e., which pass through an appropriate temporal sequence of intermediate states. Rumelhart and Norman (1985) thus propose a somewhat stronger criterion for an analogical model in addition to the two properties mentioned above:

"Every aircraft of type x has ECM gear" (e.g., abstract nodes corresponding to the set of all x-type aircraft, the set of all ECM gear, and the set of all "having" or "containing" relations; partitions of the network into components corresponding to the antecedent and consequent of the proposition, etc.). A mental model of the same fact, by contrast, might involve tokens symbolizing x-type aircraft and tokens symbolizing ECM gear associated with one another by symbols representing containment:

x-aircraft → ECM gear
 x-aircraft → ECM gear
 x-aircraft → ECM gear
 (ECM gear)

Parentheses are placed around one of the ECM gear tokens to represent the fact that some ECM gear may be present on other types of aircraft. The key feature of a mental model, according to Johnson-Laird, is the economy and naturalness of the representation it imposes.

When new information is obtained, it is not simply appended to a list of beliefs; it is added directly to the appropriate mental model. For example, on learning that "There is an x-type aircraft at y field", we get:

y-field aircraft = x-aircraft → ECM gear
 x-aircraft → ECM gear
 x-aircraft → ECM gear
 (ECM gear)

The "inference" that there is an aircraft with ECM gear at y field can now be directly read off the updated model. Thus, Johnson-Laird's primary interest in mental models is to explain features of human cognition that seem incompatible with an account of problem-solving strictly in terms of abstract reasoning, or application of general rules.

Analogical models are not necessarily constrained in the *kinds* of things they can represent, only in the way those things are represented. Thus, Johnson-Laird distinguishes between *physical* models which represent perceived objects

and relations, and *conceptual* models which represent non-perceptual relationships. It seems clear that pilot displays in future aircraft systems will involve both types of models. Pilot functions (flying the aircraft, operating sensors, planning and executing tactics) all require the formation of mental models of the aircraft in relation to the physical world of targets, sensors, weapons and environment. Equally important, however, are conceptual models, which represent non-perceived relationships such as "able to detect," "able to jam," or "able to hit."

Johnson-Laird identifies six types of physical models: relational models (a static frame containing a finite set of entities, properties, and relations); spatial models (in which all relations, both represented and representing, are spatial); temporal models (representing a sequence of events or spatial situations in time), kinematic models (a temporal model that is psychologically continuous), dynamic models (a kinematic model that incorporates causal relations) and images (a viewer-centered representation of an underlying three-dimensional spatial or kinematic model).

In addition, Johnson-Laird (1983) describes four types of conceptual model:

1. Monadic, representing assertions about individual entities, their properties, and identities between them;
2. Relational, which introduce a finite number of relations between the entities in a monadic model (such as "there are more a's than b's");
3. Meta-linguistic, which introduce semantic relationships such as "refers to," "means," "is called," etc.;
4. Set-theoretic, which includes notions of set-membership, set properties, and relations among sets.

In terms of this taxonomy, deKleer and Brown's "device topology" is a special kind of conceptual/relational model, in which tokens are related to one another by potential causal effects. Envisioning then derives a temporally

and causally related sequence of events; i.e., a "causal model" (for deKleer and Brown) - a physical/dynamic model (for Johnson-Laird).

Uncertainty. Despite this flexibility in the types of objects and relationship that can be represented, the constraints imposed by the nature of analogical models have important consequences. Most important, we think, is the difficulty that is implied in the representation of indeterminacy, whether uncertainty about facts or about values. Suppose, for example, that we know that "Base A is west of Base C" and "Base A is west of Base B." How can we represent this in an analogical model? We have two choices:

or

A	B	C
A	C	B

Our information does not specify the relationship between B and C. However, the strict requirement of isomorphism in the analogical model forces us to choose. We cannot have a model with a direct mapping to the state of affairs it represents when we do not know what that state of affairs is.

A similar difficulty arises in the representation of uncertainty about values. Suppose, for example, that a pilot is considering three tactical options. In terms of risk to own aircraft, option C is better than option B which is better than option A. But in terms of time and fuel required to execute the tactic, B is better than C which is better than A. The pilot can conclude that C is better than A (since it is superior both in terms of risk and in terms of time and fuel), and that B is better than A (since it too is superior on both dimensions). But he does not know whether B is preferable to C or C is preferable to B.

The strict requirement of isomorphism can be relaxed in various ways to represent indeterminacy (either about facts or about values), but each approach has its drawbacks:

- o Use of multiple models--e.g.,

$$\begin{pmatrix} A & B & C \\ A & C & B \end{pmatrix}$$

The problem here is the potential combinatorial explosion as new information, and new indeterminacies, are added.

- o Injection of propositional notation--e.g.,



where the arrows represent "to the west of" or "is worse than." The problem here is that the naturalness of the mental model approach is lost; inferences can no longer be directly read off the model, since the spatial relations in the model are no longer being used representationally.

- o Utilization of more imprecise models--e.g.,

A [BC]

Here isomorphism is preserved, but B and C are lumped together as a single token. This may be a viable approach, *unless* decision making requires that the relative locations or values of B and C be known.

- o Adoption of one model by assumption, with subsequent revision if necessary--e.g.,

Assume: A B C

This is perhaps the most common method. The danger, of course, is that we may lose track of (or be unaware of) our assumptions and feel an unwarranted sense of certainty.

By contrast with analogical models, normative approaches represent uncertainty by mathematically aggregating the possibilities, thus providing an abstract level of representation that corresponds to no actually realizable state of affairs. For decision making in the context of uncertainty about facts, an "expected value" is computed for each option: i.e., a weighted average of the possibilities, in which the probabilities assigned to each possible outcome

serve as the weights (cf., Raiffa, 1968). For uncertainty about values, a "multiattribute utility" score is computed; i.e., a weighted average of the scores on different evaluative dimensions, in which measures of the relative importance of differences in each dimension serve as the weights (cf., Keeney and Raiffa, 1976). Abstractions such as these can play no role in a pilot's mental models of the world since they are averages concocted for a particular occasion, not real or even possible events; hence, despite their value in decision making, they cannot be utilized effectively to increase *causal* understanding of the situation (i.e., what will happen and when).

The main weakness of mental models (their failure to represent uncertainty) is thus a by-product of their defining characteristics (the direct or analogical representation of states of affairs) and is for that reason intimately associated with their strength (the ability to generate new knowledge). In Section 2.3 we will turn to some implications of this "weakness" for the manipulation of uncertainty in unaided problem solving.

2.2.4 Hierarchical knowledge and the nature of expertise. A major variable in the performance of a combat aircraft is the level of knowledge and experience of the pilot. It is reasonable to suppose, therefore, that pilot displays for intelligent avionics should take into account and be tailored toward the level of expertise of a particular user (cf., Cohen et al, 1982; Cohen et al., 1985). In this section, we briefly consider how the knowledge structures considered above might differ between novices and experts. It turns out that the notion of hierarchical organization plays a key role.

Figure 2-3 (which may be compared with Figure 2-1) summarizes the implications of our discussion for Rasmussen's basic framework of cognitive performance. "Rule-based" performance had been replaced by a concept of stereotypical performance, which incorporates hierarchical structure and emphasizes top-down processes by means of which higher level goals and schemas may activate lower level sub-goals and schemas. Thus, as pilots accumulate experience, they may acquire more elaborate "is-a" and "is-a-part-of" knowledge structures. More extensive higher-level knowledge and top-down processes (i.e., scripts) will permit them to interpret situations more rapidly, to anticipate events, and adopt longer time horizons of planning. At the same time, more extensive

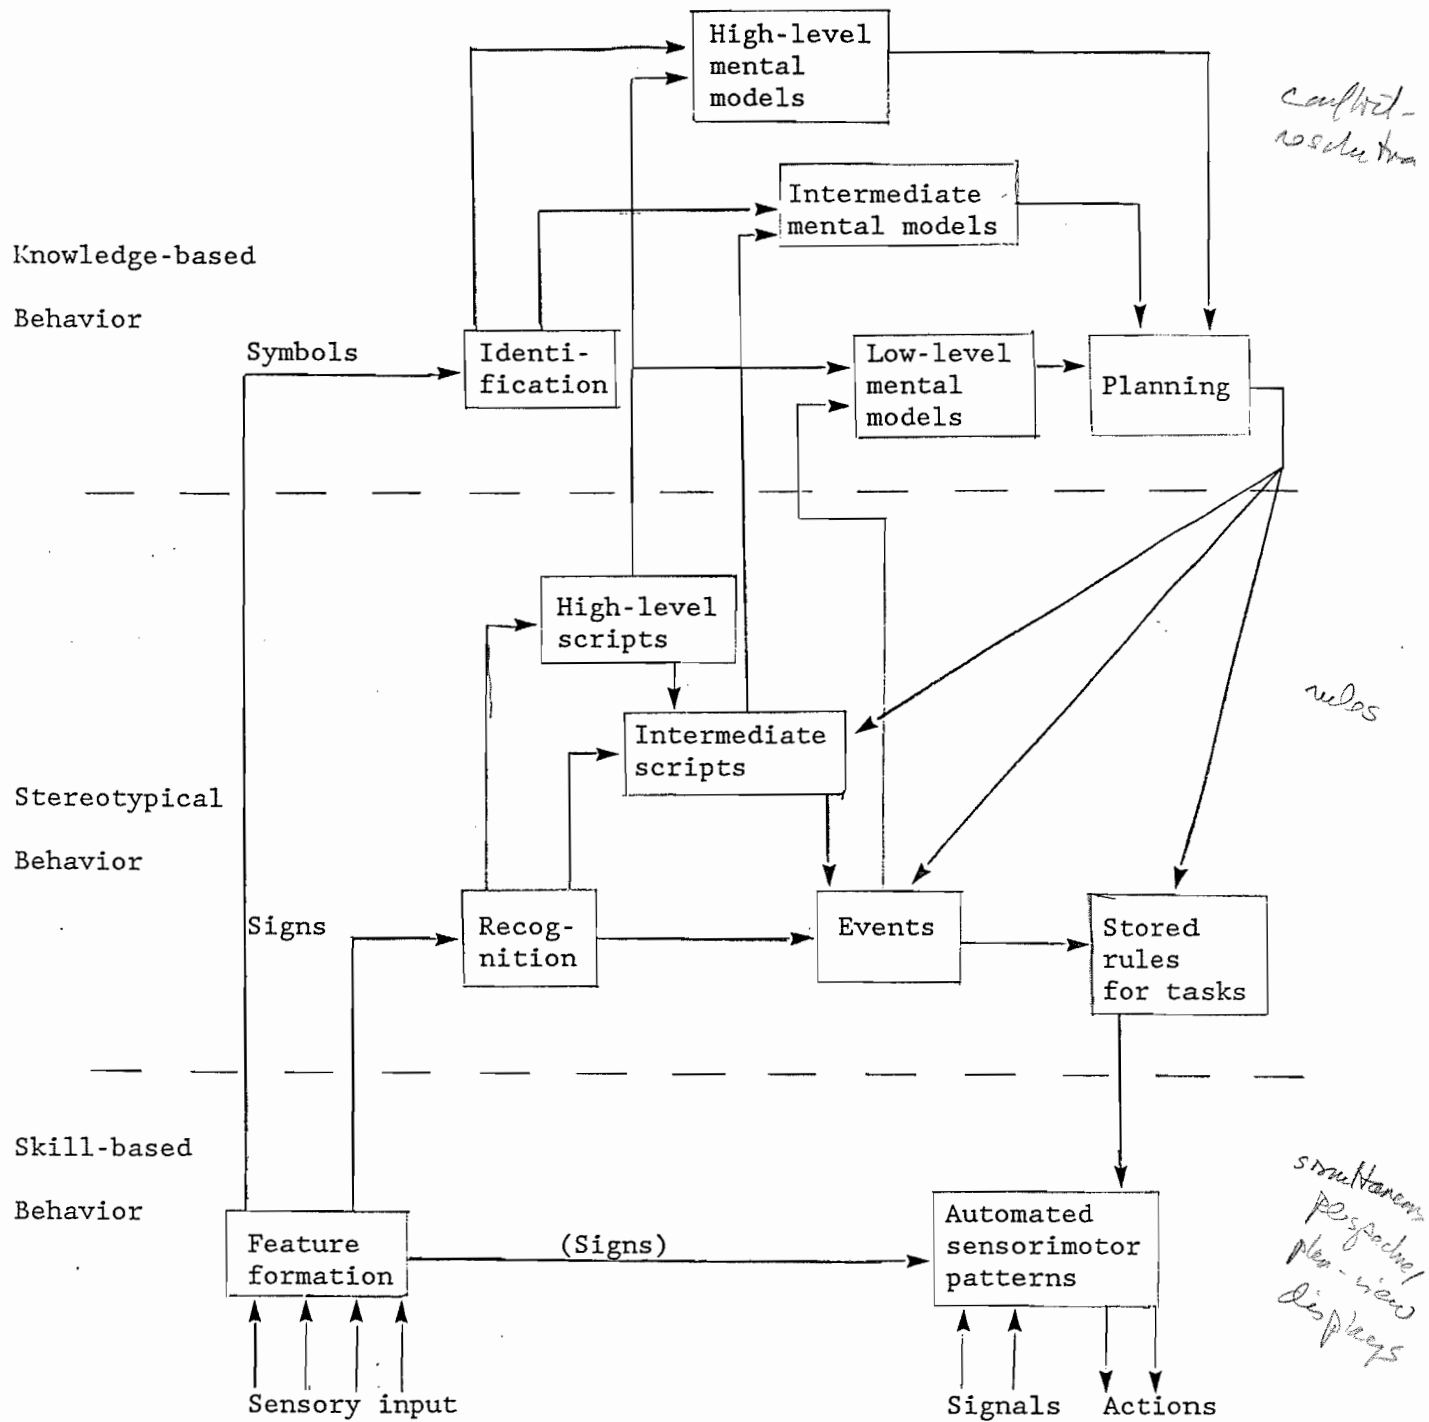


Figure 2-3: Modified Version of Rasmussen's Framework

lower level, bottom-up knowledge will permit finer discriminations among situations and more appropriate responses.

Stereotypical knowledge embodied in scripts does not represent the causal relationships underlying associations between goals and subgoals. Mental models must therefore be called upon to discover new ways of achieving goals when existing knowledge, at any level, proves inadequate (Leddo, et al., 1987). It follows that mental models themselves may differ hierarchically both on the "is-a" and "is-a-part-of" dimensions; i.e., in terms of the scope and generality of the objects which they causally relate. A pilot, for example, may take evasive action to avoid threats enroute to the target and thereby jeopardize his chances of arriving at the target by the designated time. The pilot may then use relatively low level mental models to explore alternative shorter routes or alternative faster speeds, running such models to determine if various options achieve the appropriate time over target while at the same time incurring acceptable risk. If it appears, however, that arriving at the target by the designated time with acceptable risk is not possible, higher level mental models may be utilized to decide whether ultimate mission objectives (damage the enemy, return safely) can best be achieved by continuing to the original target, aborting the mission, or seeking a secondary target instead. More experienced pilots would be expected to have developed schemas and procedures which facilitate the construction of such models, and which facilitate the selection of plausible options for testing.

A third sense in which pilot knowledge is hierarchical is represented by the classification of performance levels itself; i.e., skill-based, stereotypical, and knowledge-based. We observed above (Section 2.1) that increasing expertise is often characterized by a shifting of levels from knowledge-based to rule-based to skill-based. It is worth noting that the boundary between stereotypical performance and knowledge-based performance is not altogether clearcut. In highly unfamiliar situations, knowledge-based behavior involves the construction of a new device topology; i.e., basic components, their states, and their interconnections. However, it may be possible to deal with less novel situations by utilizing a pre-existing model and revising some of the assumptions made during "envisioning," i.e., a new temporal/causal sequence of events may need to be derived from the existing device topology.

In still more familiar contexts, it may not be necessary to re-envision the causal model; it may only be necessary to "run" an existing causal model, with new inputs representing the changed circumstances or options. Finally, in the extreme case of stereotypical performance, the stored results of previous runnings of the causal model may be retrieved to solve the present problem.

Ironically, although experienced pilots should be more skilled at building mental models, they should at the same time have less need to do so. As the number of situations with which they are familiar increases, experienced pilots have less need to call upon deeper causal analysis. Larkin et al. (1980) found that in solving physics problems, for example, sophisticated novices worked backward from the unknown through various subgoals to the given quantities and explicitly mentioned the equations used at each stage. Experts, by contrast, were faster, worked *forward* from the given to the desired quantities, and usually verbalized only numerical results rather than the equations used to derive them. These results suggest that sophisticated novices can apply generate-and-test methods (to discover ways of reducing the gap between a goal and a subgoal), but that experts have already embodied the results of such knowledge-based reasoning within stereotypical procedures. Other evidence supports the idea that expert *stereotypical* representations reflect the properties of the *mental models* from which they were derived. Chi et al. (1981) found that physics experts and novices differ in the way they sort problems by similarity. Novices categorize problems by "surface structure," i.e., superficial features such as type of apparatus, while experts rely on basic principles of physics and generic solution techniques associated with such principles. Similarly, algebra experts sort problems by solution method, while novices depend on words or objects mentioned in the problem statement (Schoenfeld and Herrmann (1982)).

In sum, there are a variety of avenues by which extended experience might affect and improve pilot performance within the framework we have described:

1. The direct accumulation of stereotypical knowledge in situations which permit (a) generalizing and aggregating lower level knowledge, and (b) refining and discriminating higher level knowledge;

2. The development of stereotypical knowledge about how to solve novel problems, i.e., increasing the ability to build and use mental models; and
3. The derivation of new stereotypical knowledge by building and running causal mental models and storing the results.

2.2.5 Behavioral decision theory. In this section we turn to a second body of research, primarily concerned with human processes of inference and choice. This work has by and large focused on errors and biases in those processes, and has been less concerned to develop explanatory models of why such errors occur. It is well beyond the scope of this report to attempt to provide such an explanation. Our objective, rather, is to suggest that the findings described above on mental models in knowledge-based and stereotypical reasoning can illuminate the nature of the errors that are observed and may provide the seeds of an eventual explanation.

Figure 2-4 provides a convenient framework for organizing the discussion of errors in reasoning. It conceptualizes the decision-making process quite generally as consisting of a specific set of cognitive tasks. First, goals or objectives must be known or identified. Secondly, current circumstances insofar as they are relevant to the achievement of goals are assessed. If a discrepancy is perceived between goals and reality, options for action may be generated. If more than one option is available a choice must be made.

This is by no means a rigid sequence: the process can be iterative (for example, revising goals, reassessing the situation, or generating new options when the choice process fails to turn up an acceptable alternative); and steps may be skipped (in particular, in stereotypical behavior when, for example, the appropriate action is known based on past experience with similar situations). Nevertheless, this framework covers the basic set of possibilities in a decision situation and, moreover, identifies the specific aspects of human performance where decision-aiding may be of use.

It is convenient to break each of these major tasks down into more specialized cognitive subtasks. For example, situation assessment consists of collecting

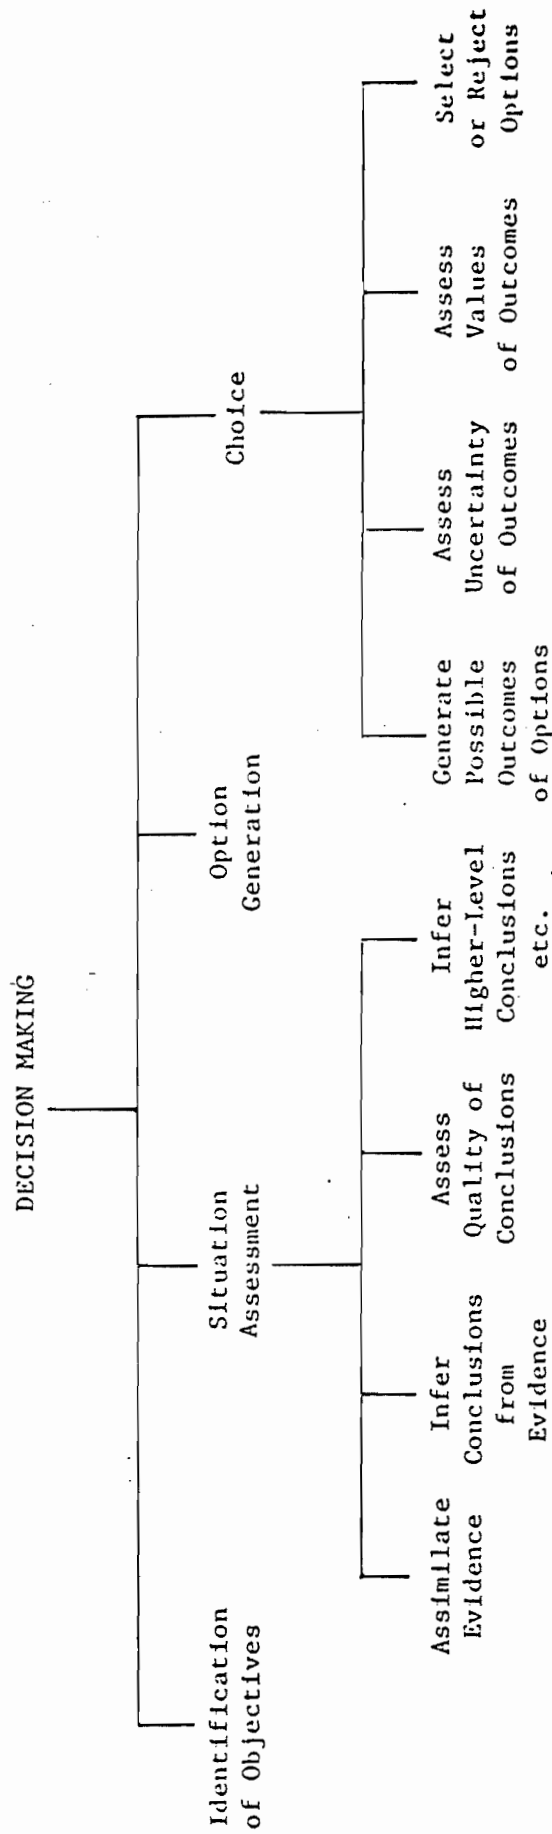


Figure 2-4: Potential Cognitive Subtasks in the Decision-Making Process.

and viewing data or evidence, deriving inferences, developing some sense of confidence in the conclusions, and continuing, perhaps, to draw further higher-level inferences. Again, the steps may be iterative, may be combined, or may be skipped altogether by some decision makers in some situations. (Note that the term "evidence" is quite relative; evidence in one process may be the highly uncertain conclusion of a prior analysis.) This decomposition of cognitive subtasks could, of course, be continued. It has been postulated that all cognitive functioning can ultimately be analyzed into a set of simple "elementary information processes" (Newell and Simon, 1972; Chase, 1978) such as selecting an input, reading the value of a variable, comparing two values, and eliminating an alternative.

Each of the cognitive subtasks identified in Figure 2-4 has been associated, at least in laboratory research, with characteristic shortcomings in reasoning. The following outline is highly incomplete and is only meant to touch on some of the issues that bear directly on the present work. Three important themes, however, should emerge: (1) Unaided decision processes employ simplifying heuristics that at best only approximate prescriptively accepted rules (e.g., Bayesian probability theory, multiattribute utility theory); (2) a typical effect of such heuristics is that awareness of uncertainty about facts or about values is suppressed; and (3) in many instances, biases are a result of (otherwise successful) efforts to utilize natural knowledge structures and processes of reasoning.

(a) *Collecting information.* Wason (1960), Einhorn (1980), and others have shown that people tend to stubbornly hold to a hypothesis generated early, avoid stringent tests of the favored hypothesis, and, in fact, seek confirming evidence. People also fail to collect evidence regarding alternative causes of an event, where more than one cause is possible (Shaklee and Fischhoff, 1982). These findings may reflect the utilization of analogical models, which are isomorphic with the states of affairs they represent and therefore fail to provide an effect representation of indeterminacy. They may also reflect the burden on short term memory and/or processing capacity of generating and manipulating more than one causal mental model.

(b) *Inferring conclusions.* A number of studies show that a statistical model of a person's judgment process can outperform (in accuracy) that person's own judgments, thus suggesting that people do not effectively utilize the information available to them in inference tasks (Dawes, 1975; Cohen, 1982). People tend to ignore later evidence that contradicts a favored, or earlier, datum and to double count redundant evidence (Schum and Martin, 1981). People commonly ignore statistical, or "base rate," data and overweight unique or problem-specific evidence, which is more readily subject to causal modeling (Kahneman and Tversky, 1972). The significance of exceptions in a series of observations is often exaggerated, i.e., treated as causally relevant rather than the result of sampling error, and, as a result, significant conclusions are overlooked (Tversky and Kahneman, 1971). These observations suggest the predominance in natural reasoning of non-statistical, causal mental models (Johnson, 1985). When people do attempt to make statistical judgments, moreover, estimates may be biased by the ease of recall (or "availability") of instances of a particular class of events in a mental sampling (Tversky and Kahneman, 1972).

(c) *Assessing quality of conclusions.* A number of studies show that people consistently overestimate their degree of certainty regarding predicted events and estimated quantities, even in areas where they are (rightfully) regarded as experts (Kadane and Lichtenstein, 1982). When inference proceeds in stages (e.g., deriving the probability of being detected by a ground radar site from information about its classification and range), people often simplify the process by acting as if conclusions at earlier stages (classification and range) were known to be true, rather than merely inferred (Schum, et al., 1973). These results also seem to reflect the difficulty of representing ambiguous states of affairs in analogical models. Similarly, the probability of a detailed hypothesis or scenario is likely to be judged higher than the probabilities for its components (Tversky and Kahneman, 1983). The latter effect may arise because additional details increase the match between the hypothesis and user mental models (Leddo et al., 1984).

(d) *Generating options.* Ingrained ways of viewing a problem (e.g., pre-existing schemas or mental models) tend to hinder the generation of novel and creative solutions. Gettys and Fisher (1979) and Gettys et al. (1981) have

shown that people often overlook important subsets of the available options or hypotheses. Moreover, people segment complex options into "natural" components (possibly based on distinct causal relationships), and treat the elements as if they were independent choices, leading to suboptimal choices (Tversky and Kahneman, 1981). There is a tendency to formulate options within a short time frame and, as a result, to overlook the cumulative risk of pursuing a course of action over a long period of time (Slovic, et al., 1978). This may reflect the difficulty of "running" mental models to simulate events far in the future, and the (related) absence of high-level aggregated schemas for novel activities of long duration. Individuals differ in the degree to which they consider future choices in current planning (Streufert and Streufert, 1981) and in the number of options they generate (Driver and Mock, 1976).

(e) *Assessing uncertainty of outcomes.* When predictions are made about the outcome of an option, there may be effects of "wishful thinking" (e.g., higher probability assessments for high utility outcomes) or overcautiousness (e.g., lower assessments for high utility outcomes) (Einhorn and Hogarth, 1984). The size of these effects may depend on the perceived uncertainty of the prediction, and may reflect a process of making assumptions to reduce indeterminacy. Perceived uncertainty in turn might depend on the degree to which available evidence matches user schemas. The "gambler's fallacy," involving distorted conceptions of randomness, may be a by-product of powerful top-down or expectancy-driven processes of pattern recognition (Lopes, 1982).

(f) *Assessing value of outcomes.* Decision makers do not typically consider all the potential outcomes of an action together. Rather, outcomes are grouped into "mental accounts" corresponding to natural objects or causal relations, and choices may depend critically on the particular grouping that is adopted (Kahneman and Tversky, 1982). Perhaps the best known research on choice behavior under risky conditions is that of Tversky and Kahneman (1981) who have shown that decisions are significantly influenced by the way the problem is framed. A key feature of this work is that people naturally represent outcomes in causally relevant terms, by the *difference* it would make relative to some reference point. Formally equivalent choice problems will be responded to differently depending on whether the outcomes are presented as

gains (e.g., lives saved relative to a worst case reference point) or losses (e.g., lives lost relative to the status quo). People tend to be risk-averse for gains and risk-seeking for losses, so that problem framing can have an important impact on choice behavior.

(g) *Selecting an option.* Choice heuristics may be adopted which reduce the amount of information which decision makers utilize. In Elimination by Aspects (Tversky, 1972), for example, attributes are considered serially in order of importance; options falling below a cut-point on an attribute are eliminated at each stage, and not considered further; and the process stops when the set of options has been reduced to the desired number. This strategy is consistent with the use of causal mental models to predict the achievement or non-achievement of a goal on each attribute. In particular, if each attribute were associated with a different mental model (for example, time to get to the target might be predicted in one model, risk to the aircraft from ground threats in another), then organizing information processing in this way minimizes the need to switch back and forth between models. The problem, of course, is that tradeoffs between goals are not considered; in this strategy, an option might be eliminated for missing a cut-point on one dimension even though it scores very high on other dimensions. Research by Lopes (1986) suggests that some decision makers compare options only in terms of their performance in the "worst case" outcome and disregard performance on other dimensions, e.g., non-worst case outcomes.

An important theme in many of these findings is that biases are a result of people's efforts to utilize natural knowledge structures and processes of reasoning. More specifically, a persuasive case can be made that biases arise from the properties of mental models: (a) the requirement of a one-to-one mapping between elements of the model and elements of the situation which they represent; (b) facilitation of the ability to "run" a single mental model, at the expense of the ability to manipulate multiple mental models simultaneously; and (c) the substitution of "inherent" relations for general rules of inference (such as in Bayesian probability theory). We have argued that all of these properties are essential for the function of mental models in generating genuinely new knowledge.

A common rationale for including humans in command and control systems is that they are more "flexible" than machines. This is presumed to mean that humans can quickly perceive new patterns or trends in a situation as it develops, and generate new hypotheses and new options that are responsive to the new conditions. As we have seen, there are important limitations to this conclusion. Nevertheless, in high-level problem solving, methods for tapping a user's knowledge will often be an important element in the success of a computerized system. For example, current (and foreseeable) artificial intelligence technology falls short of human capabilities in reasoning on multiple levels, solving novel problems (Newell, 1981), handling unanticipated types of evidence, and using concepts like causality, intention, and belief (i.e., mental models of other agents) (Buchanan, 1981; McCarthy, 1977). The results summarized above imply that techniques for exploiting such knowledge must guard against serious potential pitfalls. Thus, the design of interactive decision-aiding functions demands a precarious balancing act between encouraging, on the one hand, and modifying, on the other, a user's natural procedures for handling information.

2.3 Personalized and Prescriptive Decision Support: A Generalized Display Design Concept

In this section we draw together the threads of the previous discussions, and present a design concept for interactive displays which is based on insights both from the literature on knowledge representation and the literature on behavioral decision theory. This design concept, referred to as Personalized and Prescriptive Decision Support (PDS), permits adaptation of a system to both the decision maker (to achieve cognitively compatible displays) and the decision situation (to avoid biases), and utilizes both automatic system procedures and user choice in making the adaptation. The present discussion is based on Lehner, et al (1987); earlier descriptions are contained in Cohen et al., 1982; Cohen et al., 1985; and Cohen et al., 1986a).

This approach is in part a response to the behavioral decision making literature (discussed in Section 2.3) that suggests that human judgment and decision-making behavior are subject to a number of cognitive biases. For instance, we saw that in making choices, people often set cutoffs on separate

dimensions (and fail to consider tradeoffs), consider only some of the possible outcomes of an option, etc. The existence of cognitive biases is often used as an argument in favor of the need for decision aids. The typical approach to aiding, however, is to supplant the user's unaided method for solving the problem with a normative method, and to replace human judgment regarding the solution with the judgments provided by a normative model embedded within the aid.

By contrast, a major premise of Personalized and Prescriptive Decision Support is that user-preferred methods may have significant utility along with their flaws. Users may employ internal models that embed valuable knowledge of the problem domain accumulated over many episodes of experience. User mental models which are ideally tuned to capture complex causal relationships may, however, be quite poor at representing uncertainty or balancing tradeoffs between competing goals. Thus, cognitive biases may, in some cases, represent the downside of powerful human information-processing capabilities. Traditional decision aiding may "throw out the baby with the bath water" in forcing users to avoid biases by adopting unfamiliar modes of reasoning and representing information. The aim of Personalized and Prescriptive Decision Support is to substitute a more precise, "surgical" removal of biases--by reducing biases in the context of the decision maker's preferred approach to the problem. The goal of decision aiding is to retain the advantages of the user-preferred method (i.e., more effective exploitation of user knowledge) while producing bottom-line performance that satisfies normative constraints.

The PDS approach to the design of decision aids varies in form depending on whether adaptation to the decision maker or to the situation is primary. (1) In the former case, the primary source of initiative is the user, who determines what basic modes of representing and processing information will be used. The aid, however, provides a prescriptive back-up for this user-initiated personalization. One form of back-up involves monitoring the user's performance, comparing it to an internal normative model, and providing prompts when the user-selected strategy is likely to lead to seriously suboptimal results. (2) In the latter case, the primary source of initiative is the decision aid, which implements a normative approach to the problem. The aid, however, monitors its own performance for weaknesses (e.g., conflicting

lines of reasoning or incomplete information) and prompts the user when it concludes that the user is likely to make a significant contribution to the problem and user workload is at an acceptable level.

Figure 2-5 outlines some of the characteristics of the application that determine which of these modes should prevail. Typically, primary adaptation to the decision maker (and greater human initiative) is appropriate when there is relatively low time stress, users are relatively high-level decision makers, and the task is relatively "unstructured," i.e., options, key uncertainties, and/or dimensions of value are to some degree undefined. Primary adaptation to the situation (and greater computer initiative) is more appropriate in high-time stress, low-level, structured tasks. This distinction corresponds to the predicted dominance of knowledge-based versus stereotypical performance.

Figure 2-6 outlines the design steps involved in Personalized and Prescriptive Decision Support. The key point is to model both user strategies and a relevant normative approach. After that, the specific conditions (if any) under which a user's approach is likely to be suboptimal, according to the normative model, can be identified. At the same time, potential advantages, if any, of permitting users to deal with the problem in their preferred way are noted. The choice of a basic aiding mode depends on the features discussed above (degree of structure, level in organization, time stress), as well as on the results of the preceding steps. Thus, primary adaptation to the decision maker presupposes that there is significant value in exploiting the user's unaided approach to the problem (this is more likely to be the case in unstructured problems under low time stress). Primary adaptation to the decision maker also presupposes that any significant biases in the user's approach can be identified, and that the conditions of their occurrence can be specified.

When primary adaptation is to the decision maker, a variety of prescriptive methods may be selected to reduce the impact of biases. Specifically, as shown in Figure 2-7, the decision aid can operate in either a proactive or reactive manner, with advisory guidance that is either explicit or implicit. Guidance is proactive if it is incorporated into the design independently of any specific evidence for biased judgment on the part of a particular decision

TASK ALLOCATION INVOLVES DETERMINATION OF
BALANCE OF INITIATIVE BETWEEN HUMAN AND
COMPUTER.




Figure 2-5: Some Factors Involved in Determining Allocation
of Cognitive Tasks Between Computer and User

<u>DESIGN STEP</u>	<u>USE:</u>
I. IDENTIFY POTENTIAL <u>USER PREFERENCE</u> IN REPRESENTING KNOWLEDGE OR SOLVING PROBLEM.	COGNITIVE SCIENCE LITERATURE KNOWLEDGE ELICITATION EXPLORATORY EXPERIMENTS
II. IDENTIFY MOST APPROPRIATE NORMATIVE MODEL(S).	AL, DA, OR, ETC.
III. IDENTIFY <u>CONDITIONS OF SUB-OPTIMALITY</u> IN USER APPROACH.	MATHEMATICAL COMPARISON WITH NORMATIVE THEORY; COMPARISON OF ALTERNATIVE NORMATIVE THEORIES
IV. IDENTIFY POTENTIAL <u>ADVANTAGES</u> OF USER APPROACH.	COGNITIVE SCIENCE LITERATURE KNOWLEDGE ELICITATION EXPLORATORY EXPERIMENTS SIMULATIONS
V. CHOOSE ALLOCATION SCHEME: (A) HUMAN INITIATIVE WITH COMPUTER HELP, (B) COMPUTER INITIATIVE WITH HUMAN HELP.	FORMAL OR INFORMAL MODELS OF USER/SYSTEM PERFORMANCE
VIA. <u>DESIGN</u> AID FEATURES THAT FACILITATE BASIC USER-PREFERRED METHOD, BUT PROVIDE PROTECTION AGAINST <u>SPECIFIC</u> IDENTIFIED PITFALLS. PROTECTION MUST MESH SUFFICIENTLY WITH PREFERRED APPROACH SO THAT ITS ADVANTAGES ARE PRESERVED.	TESTS OF OVERALL SYSTEM PERFORMANCE; UTILIZATION; CONFIDENCE AND SATISFACTION OF USER - AS A FUNCTION OF SPECIFIC AID FEATURES
VIB. DESIGN AID FEATURES THAT IMPLEMENT NORMATIVE MODEL, BUT BRING USER INTO PROCESS WHERE HE CAN CONTRIBUTE. USER INPUTS MUST MESH WITH PREFERRED <u>USER</u> APPROACH TO PROBLEM, AND NOT DISRUPT HIGHER PRIORITY TASKS.	

Figure 2-6: Elements of the Personalized and Prescriptive Decision Support Approach to Decision Aid Design

PREScriptive METHODS

RECOMMENDED USER ACTION IS:

	EXPLICIT	IMPLICIT
PROACTIVE	INSTRUCTION	CHANNELING CONTEXT FAVORS MORE OPTIMAL VARIANT OF USER- PREFERRED APPROACH
REACTIVE	PROMPTING RECOMMEND ACTIONS WHICH MESH WITH BUT REMEDY SHORTCOMINGS IN USER- PREFERRED APPROACH	OUTCOME FEEDBACK

GUIDELINES FOR THE SELECTION OF A PRESCRIPTIVE METHOD:

BIAS IS A RESULT OF A WAY OF DESCRIBING OR PERCEIVING
THE PROBLEM —————> CHANNELING

BIAS IS A RESULT OF ACTIONS UNDER VOLUNTARY CONTROL OF
DECISION-MAKER —————> PROMPTING OR INSTRUCTION

OCCURRENCE OF BIAS IS NOT INEVITABLE —————> PROMPTING

BEST ACTION NOT KNOWN; LEEWAY FOR TRIAL AND
ERROR —————> OUTCOME FEEDBACK

Figure 2-7: Prescriptive Methods for Countering Potential
User Biases


maker. Guidance is reactive if it is provided in response to specific decision maker actions on a particular occasion. Explicit guidance occurs whenever the decision aid makes an explicit recommendation to the decision maker regarding his or her decision-making procedures. Implicit guidance indirectly causes modification in decision making procedures by changing the decision maker's perception of the problem or of the success of his or her current approach. *Instruction* on problem-solving procedures is thus a form of explicit, proactive guidance.

Prompting is a form of explicit, reactive guidance. Prompting occurs when the decision aid recommends a user action to remedy a possible shortcoming in results generated by the user-preferred decision process. For instance, suppose we have a decision problem where the user-preferred approach is to select a minimum risk option (Figure 2-8 gives an example of this sort.) The research reported in Lopes (1981) suggests that decision makers often select the option which does best in worst-case assumptions, while a normative approach dictates selecting the option with the highest expected value across all outcomes. One advantage of a worst case approach is that it permits the user to focus on concrete, realizable states of affairs (which can be modeled causally) as opposed to the abstract, non-realizable average or expected value. Prompting would occur when a decision aid informed the decision maker that an option existed which is slightly more risky than the minimum risk option but had a much better outcome on non-worst-case assumptions. Note that the prompt does not require the decision maker to abandon altogether his preferred mode of processing in favor of a normative approach. Rather than requiring him to think in abstract terms (i.e., to compare the expected values of each option), the prompt recommends a procedure that meshes naturally with his original approach (look only at worst outcomes), but expands it (to draw his attention to an option that does very well on better outcomes).

In the PDS approach instruction too (if it is utilized) should mesh as closely as possible with the user's natural approach, rather than impose an altogether new method (e.g., instructing such users to consider non-worst case outcomes is consistent with PDS; instructing them in expected utility theory is not) (cf., Lopes, 1982). Prompting may be preferable to instruction, however, if the potential bias does not inevitably occur whenever the strategy is used.

I) USER PREFERENCES

THERE IS EXPERIMENTAL EVIDENCE (LOPES, 1986) THAT SOME PEOPLE PREFER TO COMPARE OPTIONS IN TERMS OF THEIR ASSOCIATED WORST CASE SCENARIOS. OPTION WITH THE "LEAST BAD" WORST CASE IS SELECTED. (OTHER PEOPLE COMPARE OPTIONS IN TERMS OF ASSOCIATED BEST CASE SCENARIOS.)



II) NORMATIVE MODEL

INVOLVES ASSESSMENT OF PROBABILITIES AND VALUES OF EACH OUTCOME OF EACH OPTION, COMBINATION INTO AN EXPECTED UTILITY SCORE FOR EACH OPTION, AND SELECTION OF OPTION WITH HIGHEST SCORE.




Figure 2-8: Example of Personalized and Prescriptive Approach:
Decision Making under Uncertainty

III) POTENTIAL PITFALLS OF USER APPROACH

MAY REJECT OPTIONS WHICH ARE SLIGHTLY INFERIOR
ON WORST CASE ASSUMPTION, BUT DO BETTER IN OTHER
CIRCUMSTANCE.

IV) POTENTIAL ADVANTAGES

PERMITS A MORE INTUITIVE, LESS ABSTRACT APPROACH;
CONSISTENT WITH NEED TO ANTICIPATE AND PLAN
CONCRETELY FOR SPECIFIC SITUATIONS.

NATURAL JUSTIFICATION IN TERMS OF GUARANTEED
MINIMUM INCOME.

V) DETERMINE MODE OF AIDING

SELECT HUMAN-INITIATIVE MODE (E.G., IF THIS IS AID
FOR HIGH-LEVEL, NON TIME-STRESSED OPERATIONAL
PLANNING).

VI) AIDING APPROACH

PERSONALIZATION: UNDER UNCERTAINTY, MAKE DEFAULT
DISPLAYS CORRESPOND TO WORST CASE SITUATION.

CHANNELING: ALSO MAKE AVAILABLE DISPLAYS
CORRESPONDING TO OTHER POSSIBLE SITUATIONS, AND
TO AGGREGATED VALUES.

PROMPTING:

- PROMPT WHEN AN OPTION IS REJECTED WHICH
IS SIGNIFICANTLY BETTER ON NON-WORST
CASE ASSUMPTIONS.
- PROMPT FOR DEVELOPMENT OF CONTINGENCY
PLANS IF INFORMATION PERTAINING TO
UNCERTAINTY MIGHT BE OBTAINED LATER.

Figure 2-8: Example of Personalized and Prescriptive Approach:
Decision Making under Uncertainty (continued)

In the example above, the user-preferred (worst case) method will be significantly inferior to the normative (expected value) method only when there is an option that does very well on non-worst case assumptions but poorly in the worst case. Thus, the user need be bothered by a prompt only when it really matters.

Channeling is a form of implicit, proactive guidance. Some types of user-preferred decision strategies may be subject to very predictable biases or shortcomings. Channeling involves the tailoring of displays such that the decision maker may be less subject to these possible shortcomings. For instance, decision makers may prefer an elimination-by-aspects decision strategy (i.e., sequentially considering a series of problem aspects or factors, and rejecting options that fail to meet criteria or goals on each factor), as opposed to normative methods like multiattribute utility theory, which require explicit (and highly abstract) assessments of the relative importance of different aspects. An advantage of such a strategy, once again, is concreteness (e.g., causally modeling the achievement of specific goals on specific dimensions). Yet very good options may be inappropriately rejected because they fail to meet criteria on some factors selected for analysis, even though they perform outstandingly well on other factors. By providing displays that help users apply an elimination-by-aspects strategy while at the same time comparing options on a variety of factors, an aid may help retain the advantages of this approach while guarding against its dangers; decision makers will be less likely to reject options on the basis of a single factor. In effect, displays are designed so as to provide a context that favors the use of a more optimal variant of the user-preferred decision strategy.

Finally, providing the user with *outcome feedback* on the anticipated results of a selected decision is a form of implicit, reactive guidance. Such guidance is implicit because it leaves the user with the responsibility of discovering an emendation of his or her current procedure which will yield better performance (i.e., more satisfactory feedback). This form of guidance is thus, essentially, a matter of trial and error, and may be extremely valuable where the appropriate adaptation to a situation cannot be anticipated. In inference problems, "feedback" may consist not of "ground truth" concerning the correctness of a conclusion, but the extent of its agreement or disagree-

ment with other lines of reasoning. This type of feedback can also be based on the results of an aid-internal simulator, where simulation runs are selected to point out user-preferred vs. normative decision strategy differences.

In summary, PDS represents a form of mixed-initiative adaptation to the decision maker and decision situation. When adaptation to the decision maker is predominant, the decision aid design anticipates and provides for the possible strategies used by decision makers. The decision maker is required to have enough understanding of the decision aid and enough understanding of him- or herself to be able to select from among the alternative available decision strategies. Once a decision strategy is selected, the decision aid adapts its procedures and displays according to its internal model of the characteristics of the user-preferred strategy. In effect, the aid uses an internal model of the selected decision strategy as the major component of the model of the decision maker. This model is compared by the aid to the results of a normative model, and prompts (or other forms of guidance) are provided that help adapt the decision maker/decision aid system to the situation. These prompts are themselves influenced by the aid's model of the decision maker, and are designed to mesh closely with the user-preferred strategy. Finally, the user has the choice of determining how to respond to the offered guidance.

When adaptation to the situation is predominant, the aid utilizes a model of its own capabilities to detect potential weaknesses in its performance and a model of the decision maker's capabilities to determine when and if to prompt the user for contributions. A model of the user's preferred ways of representing information is utilized to determine the form and manner in which inputs are requested. Finally, the user may decide whether and how to respond to computer prompts.

3.0 PILOT KNOWLEDGE ELICITATION AND DESIGN OF DISPLAY CONCEPTS

In this section we review the application of personalized and prescriptive decision support to the design of interactive displays for intelligent in-flight avionic systems. As described briefly in Section 1.3, the application of that methodology proceeded in four steps:

- o Structured interviews of pilots.
- o Development of preliminary prototype displays.
- o Evaluation and comments on prototype displays by pilots.
- o Revision of prototype displays.

In principle, the last step could be followed by evaluation of the revised prototype displays, additional revision of the displays, further evaluation, and so on, until a fully satisfactory design had been developed. Within the constraints of this six-month project, however, such additional iterations were not possible.

The strategy of this section will be, first, to discuss the knowledge elicitation and evaluation methodology in somewhat more detail; we then take up several major topics in sequence: display of uncertainty, checking the validity of data sources, and hierarchical representations. Within each of these topics, we will discuss the results of each step in the application of the PDS methodology.

3.1 Method

Structured interviews. As described briefly in Section 1.3, the first stage of knowledge elicitation involved structured interviews with three pilots (as a group). The pilots were led through a typical strike scenario, in which various events were hypothesized and the pilots were asked how they would think about or act upon those events. The basic strategy of these interviews was to focus initial queries on elementary objects, events, and properties and then to gradually add complexity, e.g., multiple threats and uncertainties. Despite this structure, pilots were encouraged to talk freely about any related topics. The interviews were recorded.

During the interviews questions were asked on the following topics:

- o Approaching a single known ground threat enroute to the target
- o Approaching two known ground threats enroute to the target
- o How one thinks about own aircraft location, and the impact of location landmarks (i.e., passing a way point, getting closer to a target, passing a threat)
- o Comparing the relative danger of two threats
- o Choosing between routes which differ on various dimensions (i.e., vulnerability to different types of threats, superiority on ingress versus egress)
- o Encountering an unexpected ground threat enroute to the target under various conditions (i.e., with or without high density of surrounding threats, with or without fuel constraints, with or without limited chaff, with or without a heavy bomb load, and with or without jamming capability)
- o Encountering an unexpected air threat enroute to the target under various conditions (same as above)
- o Encountering unexpected air and ground threats simultaneously
- o Uncertainty about the location or number of ground threats under various conditions (i.e., degree of overlap with planned flight path, reliability of sources of data)
- o Uncertainty about the classification of an unexpected ground threat under various conditions (i.e., impact of uncertainty on projected flight path, reliability of sources)

- o Choice among routes which differ in risk (i.e., avoiding an uncertain but dangerous threat versus avoiding a less dangerous but known threat versus a hedging strategy)

Prototype development and evaluation. The next steps in the knowledge elicitation process involved analysis of the structured interviews, development of preliminary prototype displays, and evaluation of those displays by pilots. The displays were implemented on an IBM-PC/AT which presented the displays in the context of an illustrative ground strike scenario.

Prototype evaluation consisted of two phases, conducted individually with each of the three pilots: (a) an initial run-through of the sequence of displays in the sample scenario to familiarize the pilot with the scenario and with the basic features of the prototype system; and (2) a second run-through of the sample displays with comments and quantitative evaluations. Each of the three pilots was asked to rate twenty-four specific display features on a seven-point scale based on his experience with current cockpit equipment. 1 indicated "very good," 4 indicated "neutral," and 7 indicated "very poor." Comments were also solicited from the pilots regarding these and other display features.

A final version of prototype display system was then developed based on the pilot evaluations. Displays for that prototype system are presented and described in the Appendix in the order of the sample scenario; in the remainder of this section, however, they will be discussed in the context of specific topics to which the design methodology was applied.

3.2 Uncertainty

Structured interviews. The theory of mental models which we developed in Section 2.0 implies that decision makers in general, and pilots in particular, should experience difficulty simultaneously considering multiple possible situations, and that problem-solving efforts will be oriented towards arriving at a single acceptable, concrete (i.e., analogical) representation. This hypothesis was confirmed in the structured interviews with pilots. The

interviews revealed, however, that a variety of relatively sophisticated methods for arriving at such a representation are utilized:

- o If sensors confirm the presence of the threat but are inconclusive regarding its classification, pilots adopt a worst case assumption, i.e. they assume that the threat has maximum plausible capability against them. The rationale for this assumption is that the failure to classify the threat is itself evidence that the threat is a new system, and therefore likely to be more dangerous than previously known threats.
- o On the other hand, if available information is inadequate to confirm the existence of a threat, pilots tend to make a best case assumption, i.e., they assume that the threat is not present until more definite information is obtained. The rationale for this assumption is that actions taken to avoid the threat would almost inevitably expose the aircraft to risk from other known threats. Nevertheless, even in this situation, limited action, e.g., speeding up the aircraft, might be taken to reduce risk from the unconfirmed threat.
- o Even when the existence, location, and presence of a threat is known in advance, there may be uncertainty about its actual capabilities. Pre-briefed intelligence generally focuses on maximum capabilities, disregarding degradation during the course of combat. Pilots, on the other hand, assume that in practice all systems are subject to a significant amount of degradation; as a result they tend to apply a general discounting factor to the threat as assessed in intelligence reports.

Prototype displays. Based on the results of the structured interview, prototype displays were designed satisfying the implied constraints of mental model theory. That is, displays under conditions of uncertainty regarding threat existence, location, or classification portrayed single possible situations in preference to probabilistic averages. The particular situation depicted, however, depended on the type of uncertainty: worst case displays for classification uncertainty and best case displays for existence/location

uncertainty. As a partial safeguard against focusing exclusively on a single possibility, however, displays for other possibilities as well as an aggregated display were also made available.

Figure A-1 shows the first screen in the simulated ground strike scenario. The dotted yellow line at the bottom right represents the FEBA; the blue aircraft symbol represents own aircraft; the solid blue line represents the planned aircraft route; and the yellow "T" represents the target. Ground threats are represented by generic symbols for surface-to-air missiles, anti-air artillery, and radar. Different shades of red indicate different levels of threat to the aircraft in those regions. Figure A-7 indicates a later point in time in the scenario when new threat information has been received from an AWACS (e.g., through a JTIDS digital data link). This information suggests the possible existence of a new threat at the location indicated by the yellow lethality contour. In this scenario, however, interpretation of that data is uncertain: it could indicate the existence of a new threat; alternatively the AWACS data could represent a previously identified threat which has changed location or which was previously mislocated. Three different displays were designed to represent this situation:

- o The worst case display (Figure A-7) indicating a new threat on the planned route.
- o A best case display (Figure A-8) in which the new data are interpreted as originating from a previously identified threat, and are utilized to update the localization of that threat. (This was the default display in this scenario.)
- o An aggregated or average display (Figure A-9) in which the lethality to own aircraft at any given point is computed as a probability weighted average of the two above mentioned possibilities.

Some common features of all three displays should be noted:

- o Yellow contours are utilized to represent the receipt of new information which increases estimated danger to the aircraft. For each

display, regions are shaded in yellow when the *increase* in danger in that region, based on the new information, exceeds a pre-specified threshold (e.g., twenty percent). Estimated increments in danger to the aircraft are based on worst case and best case assumptions in displays A-7 and A-9, respectively.

- o Uncertainty is represented by the association of the red SAM symbol in the displays with a red question mark.

As shown in Figures A-10 and A-11, pilots were able to request a recommended route revision which took into account the new threat information. The recommended revision could be requested in the context of any of the three displays: i.e., a route revision based on the worst case assumption (Figure A-10), a route revision based on the best case assumption, or a route revision based on the probabilistic average (Figure A-11).

A similar set of displays was prepared to represent uncertainty about the classification of a threat. Figure A-30 represents own aircraft having passed the target and beginning the egress phase of the mission. In Figure A-31 on-board EW equipment suggests that a threat previously classified as an SA-2 may in fact be an SA-4. Figure A-31 shows the worst case assumption: that the new threat is an SA-4 (note that these threat contours are entirely fictional, and bear no relation to actual threat capabilities). Again the yellow regions indicate areas where danger to own aircraft would be increased by a given percentage on the assumption that the threat is an SA-4. Figure A-32 represents the best case assumption: i.e., that the threat is an SA-2, as previously believed. Finally, Figure A-33 represents a probabilistic average of the two possibilities. In this context, the worst case display (Figure A-31) was the default.

Evaluation. Our hypothesis, based on the theory of mental models and on the results of our structured interview, was that pilots would prefer single possibility displays (e.g., worst case or best case) to probabilistically averaged displays. In addition, we had a less strong prediction regarding which of the two single possibility displays would be preferred: worst case displays in the case of uncertainty about threat classification, and best case

displays in the case of uncertainty about threat existence/location. Finally, we proposed a prescriptive counterbalance against the likelihood that pilots would focus exclusively on single possibility displays. The menu options for the display of other possibilities and for the display of a probabilistic aggregation constitute a "channeling" device (Section 2.4) which encourages more optimal sampling of information. Since this is explicitly intended as a counterbalance to the pilot's tendency to focus on single possibilities, we did not predict *strongly* favorable responses from pilots. Nevertheless, we would expect that the display of other possibilities would conform more closely to pilot mental models, hence, be somewhat preferable to the option of viewing an aggregated display.

In the pilot evaluation of the prototype system, our main hypothesis was strongly confirmed. Pilots strongly preferred automatic presentation of displays of specific possible situations (e.g., assuming a particular threat location or threat classification) to probabilistically aggregated displays. The following table gives the quantitative evaluations of this display feature:

Presentation of specific possibilities

Existence/location uncertainty	2	1	2
Classification uncertainty	2	1	2

In this (as in all subsequent tables), the three columns correspond to the three pilots who participated in the evaluation. The pilot represented in the far right column was more senior than the other two.

However, our secondary hypothesis, regarding which automatically provided single possibility displays would be preferred under different conditions, was only partially confirmed. For uncertainty regarding threat classification, pilots did indeed prefer worst case displays. However they also preferred worst case displays when uncertainty pertained to the existence/location of the threat:

Presentation of best case

Existence/location uncertainty	6	7	6
--------------------------------	---	---	---

Presentation of worst case

Classification uncertainty	2	1	2
----------------------------	---	---	---

Pilots were then queried regarding the option of being able to see the other single possibility case. As might be expected, given its introduction as a counterbalance to the tendency to use only a single possibility, the pilots were mixed (mildly opposed, neutral, mildly favorable) in their evaluation of this option:

Presentation of other possibility

Existence/location uncertainty	5	4	3
Classification uncertainty	5	3	3

They were more mixed (mildly favorable to strongly opposed) in their evaluation of the option of seeing a probabilistic aggregation:

Presentation of average

Existence/location uncertainty	3	7	6
Classification uncertainty	3	6	5

Two of the three pilots thus felt they were more likely to use a display of the *other concrete possibility* than a display of the probabilistic average. Comments by these two pilots supported a mental model interpretation of the results. These pilots indicated that an aggregated display would be so homogenized as to be meaningless, and were confident in their own ability to extract any relevant lessons by switching back and forth between the two concrete displays.

These results suggest individual differences in the type of prescriptive channeling that most suits pilots: i.e., other single possibility displays versus probabilistic averages. The most important result, however, is that one or the other of these options was acceptable (favorable or neutral) to all the pilots, and thus might be expected to function effectively as a prescriptive counterbalance.

Pilots were strongly favorable in their evaluation of the color coded indication of increased danger due to new threat information, i.e., yellow contours:

<u>Color-coded indication of increased danger</u>	2	1	2
---	---	---	---

However, pilots also expressed a need for an additional, auditory warning in these circumstances.

Pilots were also strongly favorable in their evaluation of the system's capability of providing a recommended route revision to accommodate new threat information:

<u>Recommended route revision</u>	1	1	2
-----------------------------------	---	---	---

However they were strongly mixed when asked whether route recommendations should be provided at the pilot's request (as in our prototype) or automatically:

<u>Routes provided at pilot request</u>	1	6	5
---	---	---	---

One pilot strongly preferred that such recommendations be provided only at the pilot's request, while the other two pilots had reasonably strong preferences for the automatic provision of such recommendations. Again, the data suggest individual differences, which could perhaps be accommodated in a final system.

Prototype System Revision. These data provide support for both the personalized and prescriptive aspects of PDS (Section 2.4). The effort to tailor displays to user-preferred methods of representing knowledge and solving problems was successfully accomplished by means of the theory of mental models, according to which pilots prefer automatic presentation of single possibility displays in the context of uncertainty. The prescriptive aspect of this system guards against the tendency to focus exclusively on such a display, by providing users with the option of viewing either other single possibility displays or a probabilistic aggregation. One or the other of these two prescriptive options proved to be acceptable to all of the pilots.

Nevertheless, the choice of *which* single possibility display to present under what circumstances proved more complex than we anticipated. Additional experimentation and iterations of the prototype system would be required to fully explore this question. A plausible hypothesis, however, based on the initial structured interview as well as pilot comments during the evaluation session, is the following:

In cases of *conflict* of evidence, i.e., where there are plausible arguments on both sides, pilots consistently adopt the worst case assumption. This applies whether uncertainty pertains to location/existence or classification of a threat. On the other hand, in cases where evidence is *incomplete*, i.e., the available evidence points in one particular direction but is insufficiently reliable to substantiate that possibility, pilots have a greater tendency to adopt a best case assumption. In particular, best case assumptions will be favored if actions based on the worst case possibility are associated with known cost (i.e., increased risk from other, known threats).

In the final version of the prototype system, displays were designed to reflect this hypothesis. Thus in both of the conflict situations described above (uncertainty about location/existence and uncertainty about classification), the default display provided to the pilot represented the worst case, while the pilot had the option of viewing the best case or the aggregated display. In addition, however, we created another situation, earlier in the mission, to represent incompleteness of evidence. Thus in Figure A-2 the aircraft has received a message by electronic data link from the AWACS suggesting the possible existence of a threat on its route. Since this evidence is regarded as insufficiently reliable on its own to establish the existence of such a threat, and has not as yet been confirmed by any other data source, the system adopts a modified best case assumption. The possible existence of the threat is indicated by an empty yellow contour line and a question mark. If he wishes, the pilot may also view the worst case possibility, as shown in Figure A-3. A few moments later in this scenario, the existence of a new threat is confirmed by on-board radar. As shown in Figure A-4, when this occurs, the inference mechanism in the system regards the existence of a new threat as established and displays to the user reflect that

conclusion. In Figure A-5 the pilot has requested a recommended route revision based on the existence of such a new threat.

3.3 Validity Checking of Data Sources

In order to arrive at a single concrete representation of an ambiguous state of affairs, pilots must engage in relatively sophisticated processes of problem solving. Such processes were touched on earlier in our discussion of mental models (Section 2.2): both deKleer and Brown and Johnson-Laird focused on the use of *assumptions* to derive a concrete, analogical representation. In Section 3.1 above we confirmed that pilots engage in processes of this sort. For example, when evidence for the existence of the threat is incomplete, and avoiding the threat would incur risk, then pilots assume the threat does not exist. When there is conflicting evidence, i.e., evidence pointing in both of two directions, we saw that pilots tend to assume that the situation with greatest impact on their mission, i.e., the worst case, is true.

More active problem solving strategies are, however, available to the pilot. When evidence is incomplete, he may actively seek additional confirming data. When evidence is conflicting, he may search for an explanation of the conflict and actively seek to resolve it by revising assumptions about the sources of data. Our hypothesis regarding these more active processes is based on the theory of mental models laid out in Section 2.2 above: that pilot problem-solving strategies for dealing with incomplete or conflicting data will utilize concrete, causal models of sources of data and of the factors which might enhance or interfere with their accuracy.

Structured interviews. The structured interviews dramatically confirmed this hypothesis. While in flight over enemy territory, pilots do not simply accept pre-briefed intelligence regarding threat locations and classifications. Rather, they use such intelligence as a fallible guide in an active process of seeking additional information. In this process, the pilot continuously cross-validates information from his own sensors and from communications sources with prior expectations based on pre-briefed intelligence. When data sources do not agree, moreover, the pilot calls upon his causal understanding of the factors that affect each source in order to adjudicate the conflict.

For example, a major source of uncertainty with respect to pre-briefed intelligence is the mobility of SAM sites. Such mobility is greater close to the FEBA than it is deep within enemy territory. Therefore, other things being equal, the credibility of pre-briefed intelligence relative to other sources of data will be greater during deep penetration phases of the mission. In general, pilots attach more credibility to more recent in-flight information which is received from friendly returning aircraft, AWACS, ABCC aircraft, or own sensors. These sources, however, are also subject to error: for example, radar data may be affected by ground reflectance, weather, or electronic countermeasures. The pilot himself will often be in a position to verify, either visually or through instruments, whether any of these conditions obtain and will evaluate data sources accordingly. This continual process of re-evaluation and cross validation may not only provide a resolution of the immediate conflict, but also provides a longer term cumulative assessment of the credibility of the various information sources. For example, repeated failure to confirm RHAW scope warnings (of illumination by a threat) through other data sources may lead pilots to disregard or even turn off that piece of equipment.

In accordance with the theory of mental models, this problem-solving process is causal and qualitative rather than numerical and statistical. In the interview pilots made it clear that they did not wish to think about uncertainty in a numerical fashion.

Prototype displays. Conflict of evidence represents an anomalous (although not altogether infrequent) situation which often leads to knowledge-based reasoning on the part of the pilot. Such reasoning, and the construction and manipulation of mental models which it entails, demands considerable cognitive effort. The aim of personalized and prescriptive decision support (Section 2.5) in this context is to automate aspects of this reasoning process which can be adequately taken over by a computer, while continuing to tap the pilot's knowledge and judgment only on those occasions where he can uniquely and significantly contribute to a solution. Moreover, to maximize the pilot's contribution, displays should be designed which are compatible with his causal mental models of the data sources and which relieve some of the burden on memory and computation involved in constructing and running such models.

A set of prototype displays was developed with these objectives in mind. They support the pilot both under conditions of incomplete evidence, where he must actively search for additional data, and under conditions of conflicting evidence, where he must actively search for conditions that would causally discredit one or more of the pre-existing data sources.

Figures A-2 through A-6 illustrate the function of these displays under conditions of incomplete evidence. (As noted in Section 3.2, these particular screens were developed for the final version of the prototype system, and were not provided specifically to the pilots for evaluation.) In these screens each potential source of data regarding a threat is graphically represented by an icon. Thus on the left side of Figure A-2, from top to bottom, the folder represents pre-briefed intelligence, the aircraft stands for on-board sensors or pilot visual observation, and the lightning bolt stands for communications from air or ground stations. What a source of data has to say about a particular threat is represented by its color: a green icon means that the corresponding data source supports the best case possibility; a red icon means that the corresponding data source supports the worst case possibility; finally, a blank icon means that no reliable data has been obtained from that source. The essential idea, therefore, is to enable the pilot to see at a glance how much support there is for a particular possibility and where that support is coming from.

In Figure A-2, for example, the red lightning bolt indicates that the data link source (i.e., the AWACS) supports the existence of a new threat along the planned route; the blank icons, however, indicate that this data is not confirmed: pre-briefed intelligence and on-board sensors respectively have provided no reliable information on the presence or absence of this threat. Figure A-2 is designed to make all this information visually accessible to the pilot in an instant. In Figure A-4 the data source icon representing own aircraft sensors has turned red. This visual cue, accompanied by an auditory alert, immediately informs the pilot that the initial report of a new threat has been confirmed.

Figures A-7 through A-16 (which were provided to the pilots for evaluation) illustrate the function of this display design under conditions of conflicting evidence. In Figure A-7 data sources point to two different possibilities: either there is an unexpected surface-to-air missile site along the planned route of the aircraft, or a previously identified threat has moved or was previously mislocalized. By glancing at the iconic display, the pilot can quickly diagnose the extent and nature of the conflict among data sources. An iconic display in which all icons were green or red would indicate complete agreement. In this case, the nearly equal mix of red and green reflects extreme disagreement. The AWACS, represented by the red lightning bolt, supports the existence of the new threat; pre-briefed intelligence, represented by the green folder, supports the view that no new missile sites have been introduced into the area. Own aircraft sensor information, represented by the red and green aircraft symbol, is consistent with both possibilities. In addition, an explicit verbal indicator of "CONFLICT" is also provided.

The iconic displays do more than simply inform the pilot, in a visually immediate manner, about the current situation; they also enable him to contribute his own knowledge in an active way to the resolution of the conflict. This capability is made possible by an inference mechanism described in Cohen et al. (1986b). That inference mechanism differs in a significant way from standard normative approaches (e.g., Bayesian probability, Shaferian belief functions, or fuzzy logic) in its treatment of conflict. Rather than numerically aggregating divergent sources of information, it initiates a process of heuristic reasoning which attempts to determine and correct the cause of the conflict. It thus interprets conflict among data sources as a symptom of erroneous assumptions regarding the validity of one or more of those sources. The system attempts to resolve the conflict by selectively revising assumptions--collecting additional data to confirm or disconfirm such assumptions where possible.

Collection of additional data to resolve conflict, e.g., through deployment of on-board sensors or through communication with other ground or air stations, is determined by an automatic process which weighs the benefits against the costs of doing so. In the present displays, this data collection process has been augmented to include an interactive capability for tapping the knowledge

of the pilot. Thus if resolution of the conflict among competing data sources is significant for mission success or aircraft safety, if the pilot is likely to possess information which might help in the resolution of that conflict, and if pilot workload is at an acceptable level, then the system may query the pilot regarding factors that would potentially discredit one or more of the data sources. For example, in Figure A-13 the system has asked the pilot whether the presence of electronic countermeasures, which would invalidate the AWACS evidence, is likely. The pilot may respond to this query, ignore it, or indicate "no information." In the latter case, the system will utilize other methods for resolving the conflict, possibly including another query to the pilot.

In Figure A-14, the pilot has responded to the query by indicating that ECM affecting AWACS is indeed a problem; the icon representing the AWACS evidence has changed from red to blank; and the conflict has been resolved.

Evaluation. Two of the three pilots regarded the use of colored icons to represent agreement and disagreement among data sources favorably, while one pilot was mildly unfavorable:

<u>Icons representing conflict</u>	5	3	2
------------------------------------	---	---	---

It should be noted that the most experienced pilot was also the most favorable in his judgment of this display. A further (unscientific) observation is that approval was correlated with the order in which the pilots were exposed to the prototype system; we suspect our own skills in explaining the meaning of the iconic display improved with practice.

Two of the three pilots (although a different two) responded favorably to the use of a blank icon to represent a data source which has been discredited in the process of conflict resolution:

<u>Blank icon for discredited source</u>	3	4	3
--	---	---	---

On the other hand, pilots strongly approved the explicit indicator of conflict among data sources (i.e., the word "CONFLICT" in yellow):

Indicator of conflict

2

1

3

It should be noted that the most senior of the three pilots, who had been most favorable toward the colored iconic representation of conflict, was the least favorable toward the explicit verbal indicator.

The pilots were also favorable towards the opportunity to provide their own judgmental inputs for the conflict resolution process:

Judgmental inputs

2

2

3

Querying of the pilot by the system however was acceptable *only* on the condition (a) that the pilot was not compelled to respond, and (b) that such queries would only occur when the problem was *really* important. One pilot (the most senior) expressed an interest in the ability to directly adjust the credibility of a data source, rather than indirectly through responses to system queries.

All pilots were strongly in favor of an automated sensor management capability to guide the collection of additional data for the resolution of conflict:

Automated sensor management

2

1

2

The pilots felt that the pilot should be queried for permission to redeploy a sensor only when the pilot himself was currently utilizing the sensor to be redeployed.

Final prototype system. These results, taken as a whole, support the hypothesis that pilots deal with uncertainty by utilizing mental models of the sources of data and that displays which graphically represent what those data sources have to say can effectively support pilots in that process. We felt, however, that a more acceptable introduction to the iconic displays could be provided by a screen which was less complex than Figure A-7 or Figure A-13. This provided another motivation, in addition to those discussed in Section

3.1 above, for introducing Figures A-2 through A-6 into the final prototype system.

The ultimate objective of these displays is to provide a means whereby pilot knowledge can be effectively tapped without excessively burdening the pilot or delaying the system response. The success of these displays in that regard must eventually be evaluated in more rigorous empirical tests. Nevertheless, the pilots themselves responded quite enthusiastically to the opportunity to insert their own judgments in the conflict resolution process. In the final version of the prototype system, this capability was extended somewhat to permit the pilots to discredit a data source directly (by pointing and clicking on the relevant icon), in addition to indirectly discrediting it by responding to system queries.

3.4 Hierarchical Knowledge Representation

Pilots must of necessity think about their mission on a variety of levels. In planning, for example, they must keep in mind the overall objectives of arriving at the target with the required ordnance by the designated time and returning safely with the aircraft; a route is designed which, taken as a whole, is expected to achieve those objectives. In flight, on the other hand, the pilot's horizon of attention may expand or contract radically, depending on the circumstances. On occasions, his primary concern may be arriving at the next way-point at the appropriate time; at other times his only concern may be the immediate evasion of an active threat; on still other occasions, he may need to balance speed versus safety in replanning a significant portion of his route in the face of new information. The hypothesis to be investigated here is two-fold: (1) that displays should be appropriate to the "world" in which the pilot is currently operating, and (2) the transition from one "world" into another may be facilitated by providing displays that are (a) mutually consistent and which (b) can be continuously transformed from one into the other.

Structured interviews. Pilots think of their world from two extreme points of view, corresponding roughly to altitude. Other contrasts, which also characterize pilot knowledge representations, were mentioned during the

interviews (e.g., between planning and flying; ingress and egress). But the two extremes based on altitude were particularly significant, and the transition between them particularly difficult; a description of them will suffice to illustrate the hierarchical aspects of pilot mental models.

During a significant part of their mission pilots are performing essentially a navigation function. They are at high altitude, their geographical area of awareness is relatively large, and their temporal horizon of concern is relatively far into the future. Since they are flying above any terrain features that might be hazardous, their model is essentially two-dimensional; terrain features serve mainly as navigational cues, and their main concern is with the combined spatial/temporal goals of following a route that will avoid threats and reaching waypoints and target within a prescribed window of time. Under these circumstances, pilots rely primarily on "God's-eye" map-like displays that conform to this high-altitude, two-dimensional, large-area, long-time-horizon model.

At the other extreme, some of their time is spent flying low to avoid radar detection, maneuvering at low altitudes to evade missiles, or engaging in dog-fights with enemy aircraft. Under these conditions their geographical area of awareness is quite small, and their time horizon is of very short duration. They are intensely concerned about potentially hazardous terrain features, and their model is very much a three-dimensional one. Under these conditions pilots place a heavy reliance on direct vision outside the cockpit, and almost none on cockpit displays. Direct vision is important to them for seeing (1) missiles they are trying to evade, (2) landmarks on the bombing run, (3) terrain when it is being used for masking at low altitudes, (4) terrain that may be hazardous, and (5) air-to-air threats.

Despite the large difference between these two world models, pilots can ill afford to neglect one situation completely while operating in the other. In the interviews, they emphasized the importance of "thinking ahead of the aircraft", in the sense of mentally preparing to respond rapidly to changing situations. Thus, while flying at high altitude, they must anticipate the need to reduce altitude quickly to avoid threat tracking radar or to evade a launched missile, and mentally rehearse the actions they would take if

necessary. Similarly, after low altitude maneuvers, they must anticipate the need to make up time (either by a route change or speed change) at high altitude in order to achieve their desired time on target. One of the pilots stated that shifting from one point of view to the other took time.

Prototype displays. To facilitate the pilot's transition between the high-altitude and the low-altitude condition, a series of displays was developed to present sequential views during descent and ascent. The high-altitude display was always shown simultaneously in the upper right-hand portion. The descending transition displays (Figures A-18 to A-22) present a continuously evolving change from a high-altitude, 2-D, wide area view to a low-altitude, 3-D (perspective), small-area view. During this transition, the threat lethal contours evolve into cones shown in front of the aircraft, terrain features are shown as peaks and valleys in a head-on view, and the originally planned flight path becomes foreshortened.

During the ascending series (Figures A-23 to A-28), the reverse sequence is shown, and three features are added: (1) a recommended route for mission recovery, (2) a recommended speed for recovery of time on target (TOT), and (3) a recommended altitude for achieving the required speed with economical use of fuel. The recommended speed and altitude are also shown on the top-level high-altitude display (Figure A-28) when that final step in the sequence is shown.

The concept allows for these transition displays to be shown either before a change of altitude, to give the pilot a preview, or during the change to help him orient to the new conditions.

In the sample scenario, the descent sequence begins after the pilot views a display indicating illumination by a threat radar (Figure A-17).

Prototype evaluation. Pilots responded favorably to the simultaneous presentation of perspective and plan-view displays:

<u>Simultaneous 3-D and 2-D displays</u>	3	2	2
--	---	---	---

The most senior pilot suggested that perspective displays should be provided at all times on the heads-up display, corresponding to what the pilot would see if he were to descend to low altitude.

One pilot thought the ability to preview a continuous descent sequence was a desirable feature; the others, however, thought the transitional displays were not needed. Responses were generally the same to the display of high to low altitude transition *before* or *during* the descent:

High-to-low altitude transition

Before descent	2	6	6
During descent	2	7	6

The pilot favoring the transition display thought it would be especially valuable during night or poor visibility conditions.

None of the pilots saw value in the sequence of transition displays as a preview before ascent or as a display during ascent:

Low-to-high altitude transition

Before ascent	5	6	6
During ascent	5	6	6

However, here again the value of *simultaneous* plan-view (high-altitude) and perspectival (low-altitude) displays was noted.

Responses were highly favorable to the display of a recommended route for recovery of flight plan:

<u>Recommended recovery route</u>	2	1	3
-----------------------------------	---	---	---

With respect to the display of recommended speed to recover time on target (TOT), two pilots were highly favorable and one was neutral:

<u>Recommended recovery speed</u>	1	1	4
-----------------------------------	---	---	---

The neutral pilot thought that whether one was ahead or behind TOT would be obvious from a display of projected arrival times at check points, and that one would either speed up or dawdle, as necessary. If TOT could not be achieved, however, he would want to be informed.

Pilots were strongly favorable toward the indication of being illuminated by a threat:

Threat illumination

2 1 2

One pilot pointed out that some way of de-cluttering, or distinguishing among multiple threats in terms of priority and/or level of confidence, was needed.

Final prototype system. The pilots' evaluations confirmed the hypothesis that transitions between different cognitive "worlds" in which pilots must operate may be facilitated by simultaneous, mutually consistent displays representing those worlds. In particular, low-altitude, three-dimensional displays prior to and during descent may help pilots prepare for sudden evasive action in terrain; and high-altitude plan-view displays may help pilots regain a large-scale situation understanding prior to or during ascent. Nevertheless, pilots saw little value in sequential displays which depicted the transition between the two worlds. Other displays that supported the pilots' ability to anticipate new circumstances included recommended route and speed for recovery after a low-altitude evasive maneuver.

No changes were made to the prototype system in regard to these displays.

4.0 CONCLUSIONS

4.1 Summary of Findings from Phase I

Phase I was successful, both on theoretical and a practical level. On the one hand, some new insights into the cognitive foundations of pilot performance were obtained from a review and analysis of the cognitive science literature. On the other hand, implications of those insights for pilot displays were extracted and successfully tested. Among the more theoretical conclusions of the Phase I work are the following:

- o Pilot performance can be represented at three different levels, involving skill-based, stereotypical, and knowledge-based performance.
- o Stereotypical performance requires a characterization in terms of highly structured, hierarchical and active processes. This type of knowledge can be represented in a framework of schemas and scripts.
- o The necessary representational properties of mental models can be derived from their function of generating new knowledge. That function implies that some version of a generate-and-test process is utilized within the organism. Such an internal generate-and-test process, in turn, implies a knowledge representation in which well-understood components are "glued" together in order to observe their interaction. Such a knowledge representation is, in fact, a type of "analogical" model, in which the components correspond one-to-one with represented objects in the world, and in which conclusions are "read off" from the model itself, without the benefit of previously existing general rules or knowledge.
- o While analogical models have great strengths in supporting the ability to generate new knowledge, they are unable to represent indeterminacy or ambiguity effectively.

- o A large body of research supports the finding that unaided human problem solving is characterized by biases and fallacies, in particular in the handling of uncertainty. We argue that many, if not all, of these biases and fallacies may be explained by the human use of mental models. It thus follows that many of the weaknesses in human reasoning are inextricably intertwined with the strengths of human reasoning, i.e., the ability to use mental models to generate new knowledge.
- o An important conclusion is that there is a requirement for a design technology that both accommodates natural human knowledge structures and at the same time helps users avoid the inherent pitfalls in those structures.
- o A design methodology of this type, called personalized and prescriptive decision support, is proposed. This methodology involves modeling both user cognitive processes and representations, on the one hand, and normatively correct solutions to the problem on the other hand. These models are compared, and the potential strengths and weaknesses of the user-preferred approach are determined. Displays are designed which preserve the strengths of the user-preferred approach, i.e., which do not require users to adopt radically different "normative" techniques of problem solving. At the same time, however, these displays guard against specifically identified shortcomings in the user approach. Therefore, the end result should be performance which satisfies the constraints of the normative model, while at the same time more effectively communicating with the user and eliciting on-the-spot user knowledge.
- o Personalized and prescriptive decision support may take either of two forms. In one case, adaptation to the user's mode of problem solving is primary. The display facilitates the user's preferred approach, but monitors his performance and prompts him when his own approach is likely to lead to serious errors. In the other approach, adaptation to the situation via the normative model is primary. However, the computer monitors its own performance, and in

cases where it detects weaknesses or conflict in its own line of reasoning, where the user is judged to have potentially valuable information, and where the user's workload is at an acceptable level, the aid prompts the user for a contribution.

Pilot mental models were elicited in structured interviews in which pilots answered questions about the objects and parameters of concern to them in a series of hypothetical situations. Displays were then designed to conform to the constraints imposed by the theory of mental models. These displays were implemented in a demonstration computerized system which was then reviewed and evaluated by pilots. Preliminary conclusions and candidate display concepts include the following:

- o In cases of uncertainty about threat location or threat identity, pilots prefer displays of specific possible situations (e.g., that assume a particular threat location or type) to displays that probabilistically aggregate over the alternatives. Aggregated displays correspond to no actualizable situation, and thus may disrupt the pilot's effort to "stay ahead of the airplane" with mental models that concretely anticipate future circumstances.
- o In cases of conflicting evidence, pilots prefer situation displays which represent "worst case" as opposed to "best case" assumptions about threat location or identity, i.e., displays that depict the possibility with the greatest potential impact on the mission.
- o Nevertheless, pilots found the option of viewing a best case scenario highly acceptable. Two of the three pilots preferred the ability to compare worst and best case scenarios for themselves, rather than viewing an aggregated "average" scenario. Such options provide a counterbalance to the pilot's tendency to focus exclusively on a single possibility.
- o In cases of incomplete evidence about a new threat, pilots appeared to adopt a modified "best case" assumption, especially if taking action in regard to the new threat would itself incur risk.

- o Recommended routes for avoiding an unanticipated ground threat were strongly welcomed by pilots. As a prescriptive counterbalance to the tendency to focus on a single concrete assumption (worst case), such recommendations could be accompanied by prompts when the best response on a worst case assumption would be significantly inferior to a strategy of "hedging" against uncertainty.
- o To the extent that pilots explicitly deal with uncertainty, they utilize mental models centered on potential sources of data. Thus, pilots attempt to correlate incoming sensor reports and radio messages with prior intelligence about expected threats along a planned route; concern is aroused when these sources are in conflict. An effective mental model display, therefore, directly depicts each of the potential sources of data regarding a threat (prior intelligence, own aircraft sensors, AWACS, etc.) as an icon. The color of the icon directly encodes the impact of that source of data (green = supports best case; red = supports worst case); while the intensity of the icon directly encodes the credibility of the source of data. The pilot can thus tell at a glance the extent and nature of any conflict (if all icons are green or all are red, there is no conflict; a mix of red and green means uncertainty).
- o Pilots felt comfortable with the idea of providing their own inputs within this framework, by reducing the credibility of one or more data sources either directly or by responding to queries (e.g., about presence of countermeasures, visibility, etc.).
- o Pilots sometimes need to think simultaneously about two "worlds"-- e.g., to plan for a possible sudden descent while flying at high altitude. To facilitate this process, pilots strongly favored the simultaneous presentation of two mutually consistent displays depicting a large-area, two dimensional long-time horizon model and a narrow-area, three dimensional, short-time horizon model.

4.2 Future Directions

The principal lesson from Phase I of this research, we feel, is that a mix of cognitive science theory and empirical testing can lead to rapid progress in the development of cognitively compatible displays. The theory provides (a) a framework for understanding how pilots represent knowledge and how such representations contribute to effective performance; (b) a set of methods for designing displays that conform with the constraints of pilot internal representations; and (c) interactive techniques for counteracting the cognitive biases with which those representations are associated. Pilots themselves play a critical role in this process. Structured interviews provide an initial test of the hypotheses generated by cognitive theory and (if the hypotheses are confirmed) help us flesh out the details of the pilot's actual internal models. Review by pilots of preliminary prototype system displays provides another test of the hypotheses and further refinement of the display concepts.

Future research will continue the application of the cognitive design methodology described in this report to a wider range of pilot in-flight decision making tasks; will incorporate the resulting displays and interactive principles into a prototype real-time pilot aid; and will hopefully lead to the development of more general guidelines and methods for the design of cognitively compatible interactive displays.

Among the areas for further research are the following:

- o *Routine performance.* Research concerning scripts and schemas have as yet unexplored implications for pilot display design. To what extent should the information presented to pilots and the modes of interaction between pilot and the aide vary as a function of current goals and activities? For example, the display of threat danger may vary significantly during planning, on the ingress, during the attack, and on the egress. Similarly, interactive methods for generating and evaluating new routes and tactics may also vary as a function of where in the mission these activities occur. These displays

may also vary as a function of specific sub-goals and conditions, such as altitude, current threat density, and fuel status.

- o *Problem solving performance.* Another area of application involves decision making tasks which the pilot faces from time to time, e.g., in determining courses of action against unexpected threats or explaining unexpected events. In these contexts, displays must be provided which are compatible with the user's cognitive style, but which at the same time provide prompts or other display features which guard against decision making biases. For example, in cases of uncertainty due to conflicting evidence, we have recommended that pilots be provided with worst case displays along with the option of viewing displays that represent other possibilities. An additional protective device against potential biases might be provided in the form of prompts which warn pilots when a course of action based on the worst case assumption may be significantly inferior to actions which exploit other possibilities or which hedge against uncertainty. Another promising area of application involves choice among options which vary on multiple attributes. For example, after an evasive maneuver the pilot may be unable to arrive at the target by the designated time with the preplanned course and speed. The process of replanning involves balancing increased risk to own aircraft against the importance of the target, as well as other factors such as dependence of other aircraft on performance of the mission and/or the possible substitution of other aircraft in the mission. Displays are needed which help pilots organize and evaluate these factors, and which guard against the danger of disregarding significant information.
- o *Human computer task allocation.* Traditionally, task allocation in human-machine systems has been course-grained and inflexible: tasks are rigidly assigned to the computer or to the user according to the purported strengths of each. The display methodology described in this report opens the way to a more flexible and dynamic approach, in which the balance of initiative between human and computer shifts back and forth as a function of workload, relative expertise, and

user preferences. The key concept is that in *all* task allocation modes, user and computer complementarity is maximally exploited. Thus under circumstances when problem-solving is under the user's initiative, the computer monitors the user's decision making behavior and provides prompts when that behavior significantly violates normative constraints. Under circumstances when problem-solving is primarily under computer initiative, the computer monitors its own performance for incompleteness of evidence or conflict among data sources, and prompts the user when the user is likely to be able to make a significant contribution.

Pilot interaction with in-flight intelligent systems remains both a highly urgent and a highly promising area for the application of cognitive science display technology.

REFERENCES

- Buchanan, B.G. *Research on expert systems* (Report No. HPP-8-1). Stanford, CA: Computer Science Department, School of Humanities and Sciences, February 1981.
- Chandrasekaran, B. Towards a taxonomy of problem solving types. *The AI Magazine*, Winter/Spring 1983, 9-17.
- Chase, W. Elementary information processes. In W.K. Estes (Ed.), *Handbook of learning and cognitive processes*, Vol. 5. Hillsdale, NJ: Erlbaum, 1978.
- Chi, M., Feltovich, P., and Glaser, R. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981, 5, 121-152.
- Cohen, M.S. *Decision support for attack submarine commanders: Target range pooling and attack planning* (U) (Technical Report 82-1). Falls Church, VA: Decision Science Consortium, Inc., April 1982. Confidential.
- Cohen, M.S., Bromage, R.C., Chinnis, J.O., Jr., Payne, J.W., and Ulvila, J.W. *A personalized and prescriptive attack planning decision aid* (Technical Report 82-4). Falls Church, VA: Decision Science Consortium, Inc., July 1982.
- Cohen, M.S., Thompson, B.B., and Chinnis, J.O., Jr. *Design principles for personalized decision aiding: An application to tactical air force route planning* (Final Technical Report). Falls Church, VA: Decision Science Consortium, Inc., July 1985.
- Cohen, M.S., Laskey, K.B., and Tolcott, M.A. *A personalized and prescriptive decision aid for choice from a database of options* (Technical Report 86-1). Falls Church, VA: Decision Science Consortium, Inc., March 1986a.
- Cohen, M.S., Laskey, K.B., McIntyre, J.R., and Thompson, B.B. *An expert system framework for adaptive evidential reasoning: application to in-flight route re-planning*. (Technical Report 86-3). Falls Church, VA: Decision Science Consortium, Inc., March 1986b.
- Dawes, R.M. The mind, the model, and the task. In F. Restle et al. (Eds.), *Cognitive theory*. Hillsdale, NJ: Lawrence Erlbaum Assoc., 1975, 1, 119-130.
- de Kleer, J., and Brown, J.S. Mental models of physical mechanisms and their acquisition. In J.R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum, 1981.
- Dennett, D.C. Why the law of effect will not go away in *Brainstorms: Philosophical Essays on Mind and Psychology*, Montgomery, VT: Bradford Books, Publishers, Inc. 1978, 71-89.
- Driver, M.J. and Mock, T.J. Human information processing, decision theory style, and accounting information systems. *Accounting Review*, 1976, 50, 490-508.
- Einhorn, H.J. Learning from experience and suboptimal rules in decision making. In T.S. Wallsten (Ed.), *Cognitive processes in choice and decision behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1980.

- Einhorn, H.J., and Hogarth, R.M. *Ambiguity and uncertainty in probabilistic inference*. Chicago, IL: University of Chicago, Center for Decision Research, June 1984.
- Engel, S.E., and Granda, R.E. *Guidelines for man/display interfaces* (Tech. Rep. TR00.27200). Poughkeepsie, NY: IBM Poughkeepsie Laboratory, 1975.
- Galambos, J.A., Abelson, R.P., and Black, J.B. (Eds.) *Knowledge structures*. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers, 1986.
- Gentner, D. and Gentner, D.R. Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner and A.L. Stevens (Eds.), *Mental models*. Hillsdale, NJ: Erlbaum, 1983.
- Gettys, C.F., and Fisher, S.D. Hypothesis plausibility and hypothesis generation. *Organizational Behavior and Human Performance*, 1979, 24, 93-110.
- Gettys, C.F., Manning, C., and Casey, J.T. *An evaluation of human act generation performance* (Tech. Rep. 15-8-81). Norman, OK: University of Oklahoma, 1981.
- Johnson-Laird, P.N. *Mental models*. Cambridge, MA: Harvard University Press, 1983.
- Johnson, E.J. *Expertise and decision under uncertainty: Performance and process*. Cambridge, MA: Massachusetts Institute of Technology, February 25, 1985.
- Kadane, J.B., and Lichtenstein, S. *A subjectivist view of calibration* (Report 82-6). Eugene, OR: Decision Research, 1982.
- Kahneman, D., and Tversky, A. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 1972, 3, 430-454.
- Kahneman, D., and Tversky, A. The psychology of preferences. *Scientific American*, January 1982, 246, 160-173.
- Keeney, R.L., and Raiffa, H. *Decisions with multiple objectives: Preferences and value tradeoffs*. NY: Wiley and Sons, 1976.
- Kosslyn, S.M. *Image and mind*. Cambridge, MA: Harvard University Press, 1980.
- Lakoff, G., and Johnson, M. *Metaphors we live by*. Chicago, IL: The University of Chicago Press, 1980.
- Larkin, J., McDermott, J., Simon, D.P., and Simon, H.A. Expert and novice performance in solving physics problems. *Science*, 1980, 208, 1335-1342.
- Leddo, J., Abelson, R.P., and Gross, P.H. Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology*, 1984, 47(5), 933-943.
- Leddo, J., Mullin, T.M., and Cohen, M.S. *A user manual for eliciting and representing expert knowledge*. Falls Church, VA: Decision Science Consortium, Inc., 1987.

Lehner, P.E., Cohen, M.S., Mullin, T.M., Thompson, B.B., and Laskey, K.B. *Adaptive decision aiding* (Technical Report 87-3). Falls Church, VA: Decision Science Consortium, Inc., February 1987.

Lopes, L.L. *Averaging rules and adjustment processes: The role of averaging in inference* (T.R.). Madison, WI: Wisconsin Human Information Processing Program (WHIPP 13), 1981.

Lopes, L.L. *Procedural debiasing* (WHIPP #15). Madison, WI: Department of Psychology, University of Wisconsin, 1982.

McCarthy, J. Epistemological problems of artificial intelligence. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Cambridge, MA: Massachusetts Institute of Technology, 1977.

Metzler, J., and Shepard, R.N. Transformational studies of the internal representation of three-dimensional objects. In R. Solso (Ed.), *Theories in cognitive psychology: The Loyola Symposium*, Hillsdale, NJ: Lawrence Erlbaum Associates, 1974.

Minsky, M. The society theory of thinking. In P.H. Winston and R.H. Brown (Eds.) *Artificial Intelligence: An MIT perspective*. Cambridge, MA: The MIT Press, 1979.

Moss, R.W., Reising, J.M., and Hudson, N.R. *Automation in the cockpit: Who's in charge?* Presented at Aerospace Congress and Exposition, Long Beach, CA: October 15-18, 1984.

Newell, A. *The knowledge level* (Report No. CMU-CS-81-131). Pittsburgh, PA: Carnegie-Mellon University, 1981.

Newell, A., and Simon, H.A. *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall, 1972.

Norman, D.A., and Draper, S.W. *User centered system design: New perspectives on human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1986.

Pylyshyn, Z.W. Imagery theory: Not mysterious--just wrong. *Behavioral and Brain Sciences*, 1979, 2, 561-563.

Raiffa, H. *Decision analysis: Introductory lectures on choices under uncertainty*. Reading, MA: Addison-Wesley, 1968.

Ramsey, H.R., and Atwood, M.E. *Human factors in computer systems: A review of the literature*. Englewood, CA: Science Applications, Inc., 1979.

Rasmussen, J. *On the structure of knowledge--A morphology of mental models in a man-machine context* (Report No. M-2192). Roskilde, Denmark: Riso National Laboratory, 1979.

Rasmussen, J. Skills, rules, knowledge: Signals, signs, and symbols and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, 1983, SMC-13(3), 257-267.

- Rasmussen, J. *Information processing and human-machine interaction: An approach to cognitive engineering*. New York, NY: North Holland Publishing Co., 1986.
- Rouse, W.B., and Morris, N.M. On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin*, 1986, 100(3), 349-363.
- Rumelhart, D.E., and Norman, D.A. Representation of knowledge. In A.M. Aitkenhead and J.M. Slack (Eds.), *Issues in cognitive modeling*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1985.
- Schank, R.C. *Dynamic memory: A theory of reminding and learning in computers and people*. Cambridge MA: Cambridge University Press, 1982.
- Schank, R.C., and Abelson, R.P. *Scripts, plans, goals and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum, 1977.
- Schoenfeld, A.H., and Herrman, D.J. Problem perception and knowledge structure in expert and novice mathematical problem solvers. *Journal of Experimental Psychology: Learning, memory, and cognition*, 1982, 8, 484-494.
- Schum, D., DuCharme, W., and DePitts, K. Research on human multistage probabilistic inference processes. *Organizational Behavior and Human Performance*, 1973, 10, 318-348.
- Schum, D.A., and Martin, A.W. *Assessing the probative value of evidence in various inference structures* (Research Report 81-02). Houston, TX: Rice University, 1981.
- Shaklee, H., and Fischhoff, B. Strategies of information search in causal analysis, February 1982. Submitted to *Memory and Cognition*, in press.
- Shepard, R.N. Form, formation, and transformation of internal representations. In Solso (Ed.), *Information processing and cognition: The Loyola Symposium*, 1975, 87-122.
- Simon, H. *The science of the artificial*. Cambridge, MA: The MIT Press, 1969.
- Slovic, P., Fischhoff, B., and Lichtenstein, S. Accident probabilities and seat belt usage: A psychological perspective. *Accident Analysis and Prevention*, 1978, 10, 281-285.
- Stefik, M. Planning and meta-planning. *Artificial Intelligence*, 1981, 16(2), 141-170.
- Streufert, S., and Streufert, S.C. *Stress and the measurement of task performance. Decision making in complex tasks* (Technical Report 3). Hershey, PA: The Milton S. Hershey Medical Center, Department of Behavioral Science, 1981.
- Tversky, A. Elimination by aspects: A theory of choice. *Psychological Review*, 1972, 79(4), 281-299.
- Tversky, A., and Kahneman, D. Belief in the law of small numbers. *Psychological Bulletin*, 1971, 2, 105-110.

Tversky, A., and Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, 4, 207-232.

Tversky, A., and Kahneman, D. The framing of decisions and the psychology of choice. *Science*, 1981, 211, 453-458.

Tversky, A., and Kahneman, D. Extensional vs. intuitive reasoning: The conjunctive fallacy in probability judgment. *Psychological Review*, 1983, 90, 193-315.

Wason, P.C. On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 1960, 12, 129-140.

Winograd, T. *Understanding Natural Language*. New York: Academic Press, 1972

APPENDIX


PROTOTYPE SYSTEM DISPLAYS

This appendix contains a description of the displays included in the prototype system, presented in the order of the sample scenario. The user of the system, however, would not necessarily see these exact displays in this exact order, since in some cases what he would see would depend on his own choices. While the displays presented here do not include all those which were developed for the prototype system, they provide a representative sampling of different user menu choices and different user actions.

Figure A-1

The scenario begins with own aircraft (blue aircraft symbol) having crossed the FEBA (yellow dotted line) on a planned route (solid blue line) to a ground strike target (yellow "T"). Ground threats are represented by generic symbols for surface to air missiles, anti-air artillery, and radar. Different shades of red indicate different levels of threat to the aircraft.

Quit
Reset




Best Average Recommend Accept Descend Ascend

Figure A-2

The aircraft has now received information, via an electronic data link from an AWACS, indicating a possible threat along its planned route. Since this data is regarded by the system's inference mechanism as insufficiently reliable on its own to establish the existence of the threat, and since it has not been confirmed by other sources, the existence of the threat is not established, and a modified best case situation display is presented to the user. This consists of a yellow outline around the region where the unconfirmed threat might exist, with a question mark indicating the uncertainty. In addition, icons to the left of the display graphically indicate the status of various data sources in regard to this threat. Each icon stands for a data source: from top to bottom, the folder represents pre-briefed intelligence, the aircraft stands for on-board sensors and pilot visual observation, and the lightning bolt stands for electronic data link messages from friendly air or ground stations (e.g., through JTIDS). Red icons support the worst case assumption (existence of the new threat); green icons support the best case assumption (no new threat); and blank icons reflect inconclusive or unreliable data. Thus, the pilot can see at a glance both *how much* support there is for a particular possibility and *where* it is coming from. The red lightning bolt indicates that the data link source (i.e., the AWACS) supports the existence of the threat; however, the blank icons indicate that pre-briefed intelligence and on-board sensors respectively have provided no reliable information on the presence or absence of this threat.

Next
Quit
Reset

AMACS: SAS at 25°




Best Average Recommend Accept Descend Ascend

Figure A-3

If he wishes to, the pilot may request a worst case display. This indicates in more detail the lethality contour of the threat, *assuming* that the threat exists.

Next
Quit
Reset

AIMCS: SAS at 25°



Best Average Recommend Accept Descend Ascend

Figure A-4

At a somewhat later point in time, confirmation for the existence of a new threat is received from on-board sensors. As a result, the inference mechanism establishes the existence of the threat, and the displayed situation now corresponds to the worst case possibility. Reflecting the new information, the data source icon indicating own aircraft sensors is now displayed in red. The pilot can again see at a glance, by looking at the icons, how much support is present for a particular possibility. The yellow contours in the situation display indicate regions where danger to own aircraft has increased, by a specific percentage, on account of the new information. An auditory alert accompanies this display.




Figure A-5

The pilot may request that the system provide a recommended route to avoid the new threat. The recommended route revision is shown in purple.

Figure A-6

The pilot has indicated his acceptance of the new route. As a result, the original route plan is revised. The new threat is now shown, like other threats, in red (as opposed to the yellow contours whose purpose was to indicate new information).

Figure A-7

At this time the aircraft receives another electronic data link message from the AWACS regarding a second possible unexpected threat. The available evidence in this situation is consistent with two possibilities: there is an additional unexpected surface-to-air missile site along the planned route of the aircraft, as shown in this figure; or a previously identified surface-to-air missile site has either moved somewhat to the northeast or was previously mislocalized. The AWACS information supports the first possibility, while pre-briefed intelligence supports the second possibility (i.e., there is substantial confidence that no new sites have been introduced into the area). Own aircraft sensor information is consistent with both possibilities. In accordance with the pilot's mental model of this situation, he is automatically provided with a worst case display (i.e., Figure A-7). The icons on the left of the screen graphically indicate the directions in which each data source is pointing. Thus, the pilot can tell at a glance whether data sources are in agreement (all green or all red) or are in conflict, as in this case. In addition, an explicit "CONFLICT" indicator is provided above the icons.

Next
Quit
Reset
Game
New
Help
Quit

AIMCS: 246 AT 300°

ECM? No Info Unlikely Possible Probable
Best Average Recommend Accept Descend Ascend

Figure A-8

If he wishes, the pilot may examine best case possibilities as well. In this display, the new information from the AWACS is interpreted on the assumption that it represents a moved or mislocalized, but previously known, threat.




Figure A-9

The pilot also has the option of examining an average, or probabilistically aggregated, display. In this display, the danger at any given point is the weighted average of the danger on the worst case possibility and on the best case probability, with weights corresponding to the probabilities of those two situations.




Figure A-10

The pilot may request a recommended route revision based on the new information. Such a revision may be requested in the context of the worst case display, the best case display, or the average display. This figure shows the recommended route on the worst case assumption.

Figure A-11

This figure shows the recommended route in the context of the probabilistically aggregated display.

Figure A-12

Here the pilot has indicated acceptance of the route based on the probabilistically averaged display. The recommended revision is incorporated into the previously planned route, and the yellow contours (indicating unexpected information) are replaced by the standard red contours. The question mark remains to indicate continuing uncertainty regarding the existence and/or location of the threat.

Figure A-13

In the case of disagreement among sources of data, the system provides the opportunity to resolve the conflict by discounting one or more of the sources. The inference mechanism automatically examines potential causes of the conflict, i.e., assumptions upon which one or more of the conflicting data sources depend for their credibility. For example, radar data may be affected by ground reflectance, weather, or electronic countermeasures. If the system can automatically resolve the conflict, it does so (by additional data collection or data analysis). If it cannot, and if the conflict is significant for mission success or aircraft safety, the system queries the user regarding factors that would potentially discredit one or more of the sources. Thus, in Figure A-13 the system has asked the pilot whether the presence of ECM, which would invalidate the AWACS evidence, is likely. The pilot is free to respond to this query, ignore it, or indicate "no information." If he indicates the latter, the system may produce an additional query.




Figure A-14

In this figure, the pilot has responded to the query by indicating that ECM affecting the AWACS is probable; the AWACS evidence has been discounted, as shown by the blank lightning bolt icon; and the conflict has been resolved by the inference mechanism in favor of the best case possibility. The pilot could have indicated his lack of confidence in a data source more directly simply by pointing and clicking at the icon representing that data source. When he does so, the data source is discounted (i.e., the icon becomes blank), and the conflict is resolved in the appropriate direction.




Figure A-15

After resolution of the conflict, the pilot requests a recommended route in the context of the best case possibility. The recommended route is shown in purple.




Figure A-16

The user has indicated his acceptance of the recommended route.




Figure A-17


Later in the scenario, on the way to the target, onboard sensors indicate that the aircraft has been illuminated by a surface-to-air threat. The yellow arrow from the threat to the aircraft represents the increased danger in this situation. An auditory alert is also provided.

Next

Quit

Reset

Detect Illumination by Threat 8



Worst

Best

Average Recommend

Accept

Descend

Ascend

Figure A-18

As a result of the threat illumination, the pilot decides to descend to lower altitude to exploit terrain. Such a descent involves a rapid alteration in the pilot's viewpoint: from a large-scale, two-dimensional plan-view to a narrow, three-dimensional perspectival view. To facilitate this transition, the system presents a sequence of views which anticipates what the pilot will see on his descent. Figure A-18 shows the aircraft in the initial portion of the descent. A plan-view situation display is shown simultaneously in the upper right.




Figure A-19

The sequence of displays corresponding to the descent continues. As the "point of view" of the display descends, it also begins to look ahead rather than down. As this happens, features of the display evolve in a continuous manner: i.e., threat lethal contours become cones shown in front of the aircraft, terrain features are shown as peaks and valleys, and the planned aircraft route becomes foreshortened.




Figure A-20

The descent sequence continues.




Figure A-21

The descent sequence continues.




Figure A-22

This is the final display in the descent sequence, showing the aircraft at its lowest planned altitude.




Figure A-23

During or prior to the ascent back to standard altitude, the system provides a corresponding sequence of displays which shows the aircraft on the ascent. Again, display features evolve continuously during the transition, and the large-scale plan-view situation display continues to be shown simultaneously.




Figure A-24

After the ascent, the pilot's objective is to recover his original flight plan to the target. Thus, this display in the ascent sequence provides a recommended route, speed, and altitude for recovery.




Figure A-25

This is the next display in the ascent sequence. The "point of view" of the display begins to look down (as opposed to forward) as the aircraft increases in altitude.




Figure A-26

This is the next display in the ascent sequence.




Figure A-27

This is the next display in the ascent sequence.





Figure A-28

The plan-view situation display is resumed, showing a recommended speed and altitude to reach the target by the designated time at minimum risk.

Next
Quit
Reset



Work

Best

Average

Recommend

Accept

Descend


Ascend

Figure A-29

A short while later, the pilot is again illuminated by a surface-to-air threat. The "X" over the yellow arrow indicates that on-board electronic countermeasures have effectively negated this threat.

Red
Quit
Reset

Detect Illumination by SA-3. Countermeasures Effective




Worst Best Average Recommend Accept Descend Ascend

Figure A-30

The pilot has successfully struck the target and is entering the egress portion of the route.

Next
Quit
Reset



Words: Best Average Recommend Accept Descend Ascend

Figure A-31

At this time, the pilot again receives unexpected information, this time pertaining to the classification of a threat. On-board EW equipment indicates that a surface-to-air site near the planned egress route may be an SA-4 as opposed to an anticipated SA-2. (Note that these displays use fictional parameters for threat capabilities. The SA-4 is thus regarded as more capable than the SA-2.) Figure A-31 shows the worst case situation which is automatically provided to the pilot, i.e., classification as an SA-4. Yellow contours indicate areas where the increased danger to the aircraft due to new information has exceeded a certain threshold. The icons on the left indicate what various data sources are saying about threat classification: i.e., green indicates support for classification as an SA-2 (best case), and red indicates support for classification as an SA-4 (worst case). The mix of red and green in the icon display thus indicates to the pilot at a glance that significant conflict exists regarding this threat. An explicit "CONFLICT" indicator is also provided, above the icons.



Figure A-32

If he wishes, the pilot may also examine the situation under best case assumption, i.e., assuming that the threat is an SA-2. Since this assumption corresponds to the prior expectation, no yellow contours are shown on this display.





Figure A-33

The pilot also has the option of viewing a probalistically aggregated, or average, display. In this display the danger at any given point is a weighted average of the dangers on each of the two possibilities, where the weights correspond to their probabilities.

DATA: Int Signal

QUI

Reset



ELINT?

No info

ECM Unlikely

ECM Possible

ECM Probable

750

Best

Average

Recommend

Accept

Descend

Ascend

Figure A-34

The pilot has requested a recommended route to accommodate the new information, based on the probabilistically aggregated display. The recommended route is shown in purple.

Figure A-35


The pilot has accepted the recommended route. As a result, the route revision has been incorporated into the preplanned route, yellow contours have disappeared, and uncertainty continues to be acknowledged by the presence of the question mark.

A-71

Figure A-36

The pilot continues the egress towards the FEBA.

Next
Quit
Reset



Worst Best Average Recommend Accept Descend Ascend